Microsoft®

Microsoft® Exchange 2000 Server Resource Kit

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

Copyright © 2000 by Microsoft Corporation

Library of Congress Cataloging-in-Publication Data
Microsoft Exchange 2000 Server Resource Kit / Microsoft Corporation.
p. cm.
Includes index.
ISBN 0-7356-1017-7
1 Microsoft Exchange Server (Computer file) 2 Client/server computing. I. Microsoft Corporation.

QA76.9.C55 M5285 2000
005.7'13769--dc21 00-055910

Printed and bound in the United States of America.

6 7 8 9 QWT 7 6 5 4 3 2

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 Prevs International directly at fax (425) 936-7329 Visit our Web site at www microvoft com/mspress Send comments to *rkinput@microsoft com*

Active Directory, ActiveX, FrontPage, IntelliMirror, JScript, Microsoft, Microsoft Press, MSDN, NetMeeting, Outlook, Visual Basic, Visual C++, Visual Studio, Win32, Windows, and Windows NT 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.

Unless otherwise noted, the example companies, organizations, products, people, and events depicted herein are fictitious. No association with any real company, organization, product, person, or event is intended or should be inferred.

Acquisitions Editor: Juliana Aldous
Project Editor: Maureen Williams Zimmerman

Part No. 097-0007796

We dedicate this volume to the authors, contributors, and reviewers who generously gave their time, experience, and knowledge to make this book a reality.

Group Program Manager Allan Risk

User Education Manager Walden Barcus

Program and Project Managers Matthias Leibmann, Sharon Farrar, Dylan Miller, Peter Van Niman, Valerie Wright

Authors Andy Koopmans, Dylan Miller, Phil Embleton, Melissa Simmons, John Speare, John Fisher, Robert Dring, Bob Hunt, Tom Richer, Paul Sebben, Patrick McFarland, Peter Nilsson, Michael Aday, Will Martin, Andrew Holmes, Ziad Chbeir, William Riedell, Ramon Infante, Christophe Besançon, Christophe Leroux, Jens Trier Rasmussen, Martyn Davis, Cherif Djerboua, Daniel Martin, Greg Dodge, Mark Garcia, Dan Bloch, Sasha Frljanic, David Hitchen, Jung-Uh Yang, Stanley Lum, Jim Pillar, Marc Stanton, Sybil Wood, Walden Barcus, Markus Vilcinskas, Tony Soper, Matthias Leibmann, Carl Solazzo

Lead Editor Frank Delia

Editors Todd Young, Susan Bradley, Amy Stockett, Pamela Miller, Sharon Farrar. Valerie Wright

Contributors Pierre Bijaoui, Luc Clement, Steve Peschka, Andreas Essing, Paul Turner, Randy Treit

Technical Reviewers Marc Stanton, Glen Anderson, Ken Ewert, Geeman Yip, John Kenerson, Peter Waxman, David Lemson, Paul Bowden, Tony Soper, Erik Ashby, Tim Kiesow, Mitch Gray, Nat Ballou, Kevin Kaufman, Mark Wistrom, Perry Clarke, Matt Gossage, Patrick Walters, Bogdan-George Pintea, Luc Clement, John Goodacre, Dalen Abraham, Rick Ryan, Chris Vandenberg, Andres Sanabria, Greg Baribault, Brian Trenbeath, Mike Gahrns, Karim Batthish, David Trulli, Andras Luther, Maureen Tracy Venti, Greg Veith, Jason Mayans, Phillip Hupf, Barry Steinglass, Linda Kadowaki, Jeff Wilkes, Margaret Li, Tom McCormick, Robert Edwards, Andreas Essing, Seigfried Jagott, Jeffrey Kempenich, Robby Voet, Adam Dare, Robert Osborne, Jeff Bachmeier, Vanitha Prabhakaran, Michio Nikaido, Marc Lauzon

Tools Dylan Miller, Kumar Rallabhandi, Mike Burke, Marcelo Calbucci, Scott Briggs, Barry Steinglass, Tom Larson, Les Thaler, Michael Patten, John Kozell, Jeff Wilkes, Kali Buhariwalla, William Finkle, Chris Falter, Ewan Dalton, Karl Froelich, Je Seog Park, John Gilbert, John Kenerson

Design, Illustration, and Production David Vican, Kristie Smith, Rick Achberger, David Hose, Buck Guderian

Project Editor Maureen Zimmerman

Table of Contents

Introduction vii

Information Roadmap xi

Enterprise Deployment Guide Table of Contents xvii

PART 1 Exchange 2000 Project Planning 1

CHAPTER 1 What's New 3

CHAPTER 2 Building the Project Plan 19

PART 2 Planning for Exchange 2000 and Active Directory 37

CHAPTER 3 The Exchange 2000 Environment 39

CHAPTER 4 Active Directory Design 63

CHAPTER 5 Active Directory Integration and Replication 109

CHAPTER 6 Deployment Strategies 149

PART 3 Prototyping Exchange 2000 175

CHAPTER 7 Setting Up a Test Environment 177

CHAPTER 8 Piloting Exchange 2000 193

CHAPTER 9 Preparing a New Environment 201

CHAPTER 10 Preparing an Existing Environment 245

PART 4 Basic Deployment Planning 279

CHAPTER 11 Administration and Maintenance 281

CHAPTER 12 Server Design for Backup and Restore 305

CHAPTER 13 Virus Protection 319

CHAPTER 14 Server Availability 333

CHAPTER 15 Server Sizing 351

CHAPTER 16 Message Routing 373

CHAPTER 17 Backbone Configuration and Tuning 383

CHAPTER 18 External Connectivity 403

PART 5 Advanced Deployment Planning 429
CHAPTER 19 Chat and Instant Messaging Services 431
CHAPTER 20 Inter-Organization Replication and Directory Synchronization 449
CHAPTER 21 Branch Office Scenarios 463
CHAPTER 22 Corporate Backbone Scenario 487
CHAPTER 23 Hosted Service Environments 505
CHAPTER 24 Security Sensitive Environments 519
CHAPTER 25 Outlook Web Access 557
APPENDIX A Client Network Traffic Analysis 573
Resource Guide Table of Contents 639
CHAPTER 26 Exchange 2000 Architecture 659
CHAPTER 27 Application Development 707
CHAPTER 28 Backup and Restore 759
CHAPTER 29 Monitoring and Maintaining 803
CHAPTER 30 Security 827
CHAPTER 31 Optimizing Exchange 2000 871
CHAPTER 32 Real-Time Collaboration 925
CHAPTER 33 Troubleshooting 947
APPENDIX B Ports and Protocols 973
Glossary 977
Index 1001

Introduction

The *Microsoft Exchange* 2000 *Sewer Resource* **Kit** contains two separate guides: the Enterprise Deployment Guide and the Resource Guide.

The Enterprise Deployment Guide contains information to help you deploy Microsoft Exchange 2000 Server in various environments. It explains upgrading earlier versions of Exchange, migrating to Exchange from other messaging systems, creating intra-company and inter-company Exchange organizations, deploying real-time collaboration services, and integrating Exchange with the Active Directory directory service. The Enterprise Deployment Guide also contains server sizing recommendations, network routing considerations, backup and restore information, project planning guidelines, and practical planning strategies.

The Resource Guide contains in-depth and advanced technical information about optimizing your Exchange organization, backing up your servers, preventing and recovering from disasters, and designing your Exchange architecture. This content includes information from the Exchange product development team and participants in the Exchange Rapid Deployment Program.

The *Microsoft Exchange* 2000 *Server Resource Kit* also includes a companion CD that contains an online version of this book, tools, sample applications, and templates.

Enterprise Deployment Guide

The Enterprise Deployment Guide contains a broad collection of advice, recommendations, and practical information that has been developed by experts from Microsoft Consulting Services and partners who are solving real deployment challenges in the field.

As you read the Enterprise Deployment Guide and implement suggested strategies, it is important to consider how the authors developed the content for this guide. The content is based on the functionality of pre-release versions of Exchange 2000. Therefore, you are reading the first available deployment information based on real-world experience with Exchange 2000. Program managers and software testing engineers from the Exchange product team have reviewed and provided feedback on each chapter. Nonetheless, remember that some of the findings from the authors of this guide and feedback from the Exchange Rapid Deployment Program might have resulted in improvements in Exchange 2000 after this book was released for publication. Thus, it is important to regularly check for updated deployment information. For the most current information, visit Exchange Up-To-Date on the Exchange Web site at http://www.microsoft.com/exchange.

Because Exchange 2000 and Microsoft Windows 2000 Server include significant improvements in functionality, no single volume can contain a comprehensive treatment for all deployment, upgrade, and migration topics. Deploying Active Directory and Exchange 2000 in a large enterprise requires careful planning, especially if you are upgrading or migrating from an existing messaging system. Good planning produces more manageable, scalable, and reliable messaging solutions. Use the Enterprise Deployment Guide as one of many sources for Exchange and Active Directory deployment information.

The content in the Enterprise Deployment Guide exists primarily for enterprise administrators to plan and execute Exchange deployment projects across large geographical areas. Administrators for smaller or less complex deployment projects can also benefit from the advanced deployment information in this guide. This information is useful for any administrator who wants to design a messaging system that does not restrict future growth.

Different professionals have written each chapter in the Enterprise Deployment Guide based on their experiences with pre-release versions of Exchange 2000. Thus, each chapter presents a contained solution rather than one part of a comprehensive narrative. Although reading this book from beginning to end is recommended, you do not need to do so to solve specific challenges. Rather, you can read those chapters that pertain to your situation and briefly scan other chapters for information that might be useful in the future.

Resource Guide

The Enterprise Deployment Guide provides strategies and knowledge from the field. The Resource Guide explains how Exchange 2000 works. The Resource Guide contains the first available advanced information for Exchange 2000. It extends the Exchange 2000 Server product documentation.

Administrators who want to learn how Exchange 2000 works will benefit from the information in this guide, including topics about the architecture of Exchange and steps for troubleshooting. This guide contains information about Exchange and Windows 2000 features and concepts. Microsoft recommends you read the *Microsoft* Windows 2000 *Server Resource Kit* to help you understand the integration between the two products.

The Resource Guide explains Exchange internal processes and provides procedures for backing up your servers, preventing disasters, securing servers from attacks, keeping users' data private, optimizing your Exchange 2000 and Windows 2000 configuration, and tuning your servers. You also learn how to troubleshoot Exchange 2000 issues and customize your Exchange organization with technologies such as Collaboration Data Objects (CDO).

This book is written for enterprise administrators and assumes that you know about Exchange concepts that are presented in the Exchange 2000 documentation. For the most current information about Microsoft Exchange 2000 Server. visit Exchange Up-To-Date on the Exchange Web site at http://www.microsoft.com/exchange.

Resource Kit Compact Disc

The *Microsoft Exchange* 2000 *Sewer Resource Kit* companion *CD* includes a variety of tools and resources to help you work efficiently with Exchange 2000.

> **Note** The tools on the CD are designed and tested for the U.S. version of Exchange 2000 and Windows 2000. These programs can cause unpredictable results. The documentation on the CD, including the HTML Help version of this book, is in English.

Exchange 2000 Server Resource Kit Online Book contains everything that is in the printed book in an online, searchable format.

Exchange 2000 Server Deployment Planning Templates developed by Microsoft consultants are for you to use and customize during the planning process.

Exchange 2000 Server Tools and Tools Help were developed by Microsoft. These tools are unsupported and provided in their current form without warranty of any kind. The tools on the companion CD are introduced and documented in an HTML Help file (Exchtool.chm) that is in the Help directory on the CD.

Resource Kit Support Policy

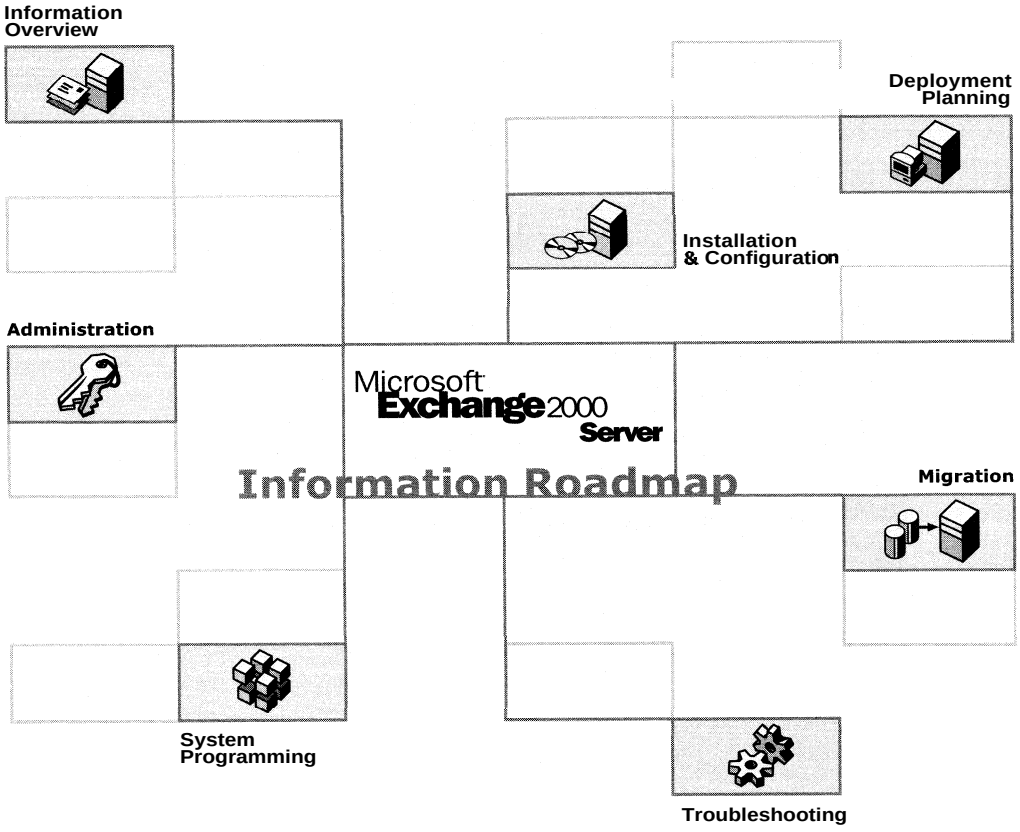
The software supplied in the *Microsoft Exchange* 2000 *Sewer Resource Kit* is not supported. Microsoft does not guarantee the performance of the tools, response times for answering questions. or bug fixes for the tools. However, Microsoft does provide a way for customers who purchase the *Microsoft Exchange* 2000 *Server Resource Kit* to report bugs and receive possible fixes for their issues. You can do this by sending e-mail to rkinput@rnicrosoft.com. This e-mail address is only for issues related to *Microsoft Exchange* 2000 *Sewer Resource Kit.* For issues related to the Microsoft Exchange 2000 Server product, see the support information included with the product.

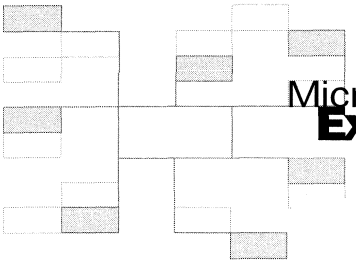
Document Conventions

The following style conventions and terminology are used throughout this book.

Element	Meaning
CAPITAL LETTERS	Acronyms, abbreviations, names of certain commands, and names of keys on the keyboard.
Initial Capitals	Names of applications, programs, files, servers, and windows; and directory names and paths.
Bold font	Menus and menu commands, command buttons, tab and dialog box titles and options, command-line commands and options, and portions of syntax that you must type exactly as shown.
Italic font	Information you provide, terms that this book introduces, book titles, and emphasis.
`Monospace font`	Examples, sample command lines, program code, and program output.

Reader Alert	Meaning
Tip	Alerts you to supplementary information that is not essential to the completion of the task at hand.
Note	Alerts you to supplementary information.
Important	Alerts you to supplementary information that is essential to the completion of a task.
Caution	Alerts you to possible data loss, breaches of security, or other more serious problems.
Warning	Alerts you that failure to take or avoid a specific action might result in physical harm to you or the hardware.


Information Overview
Deployment Planning
Installation & Configuration
Administration
Microsoft Exchange 2000 Server
Information Roadmap
Migration
System Programming
Troubleshooting

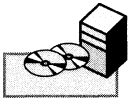


Microsoft Exchange 2000 Server

Information Roadmap

The Information Roadmap is your guide to resources for learning and running Microsoft Exchange 2000 Server. Resources are described in "Information Overview," later in this section. They include information provided with Exchange 2000 Server, information available on the Internet, and books available through Microsoft Press.

Installation & Configuration



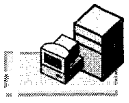
- Planning & Installation
- Resource Kit
- Administrator Help
- Exchange Web site on microsoft.com
- Microsoft Product Support Services
- TechNet

Administration



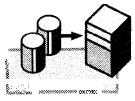
- Planning & Installation
- Resource Kit
- Exchange Web site on microsoft.com
- Microsoft Product Support Services

Deployment Planning



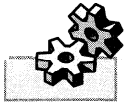
- Administrator Help
- Resource Kit
- Exchange Web site on microsoft.com
- Microsoft Product Support Services
- Exchange 2000 Newsgroups
- TechNet

Migration



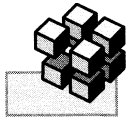
- Planning & Installation
- Resource Kit
- Microsoft Product Support Services
- Exchange Web site on rnicrosoft.com

Troubleshooting

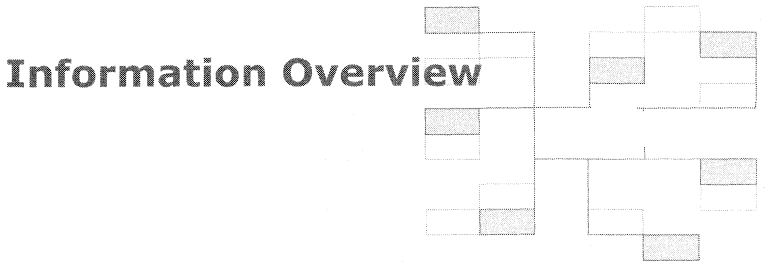


- Administrator Help
  Resource Kit
- Microsoft Product Support Services
  Exchange Error and Event Message Reference

System Programming



- Microsoft Exchange Server Developer Center
- Exchange 2000 Server Software Developer Kit (SDK)
- MSDN Online Support Newsgroups
- Microsoft Product Support Services
- Programming Collaborative Web Applications with Microsoft Exchange 2000 Server
- Programming Microsoft Outlook and Microsoft Exchange, Second Edition



- **Planning & Installation**

  This guide includes basic information about installation, product use, and features. Provided with Exchange 2000 as a printed book and as online documentation.

- **Resource Kit**

  Enterprise Deployment Guide for upgrade, migration, and integration information and Resource Guide for in-depth technical information. Available through Microsoft Press. Visit http://mspress.microsoft.com/.

- **Administrator Help**

  In the Exchange 2000 online Help, the How To... topics help you get things done, and Concepts topics explain Exchange components.

- **Exchange Web site on microsoft.com**

  Resources and technical information regarding Exchange 2000 deployment, collaboration, and administration. Visit http://www.microsoft.com/exchange/.

- **Microsoft Product Support Services**

  Web-based resource for the latest service packs, fixes, white papers, and FAQs, plus a searchable knowledge-base containing technical support information and self-help tools. Visit http://support.microsoft.com/support/exchange/.

- **Exchange 2000 Newsgroups**

  Web-based forums where you can interact with Microsoft staff and seasoned Exchange users. Visit http://www.microsoft.com/exchange/support/.

- **MSDN Online Support Newsgroups**

  Web-based forums for sharing information on developing with Microsoft products, including Exchange and Microsoft Outlook. Visit http://msdnonline.one.microsoft.com/default.asp.

- **TechNet**

  Web-based resource to help you deploy, maintain, and support Exchange. Visit http://www.microsoft.corn/technet/exchange/.

- **Exchange Error and Event Message Reference**

  Web-based source of supplemental information to explain and resolve Exchange 2000 error messages and event messages. Visit http://www.microsoft.com/exchange/support/help.htm.

- **Microsoft Exchange Server Developer Center**

  Web-based source of information about creating applications using Exchange 2000 and related technologies. Visit http://msdn.microsoft.com/exchange/.

- **Microsoft Exchange 2000 Server Software Developer Kit (SDK)**

  Web-based source of information about creating applications for messaging, communication, and collaboration using the new development platform provided by the Web Storage System and Exchange 2000 Server. Visit http://msdn.microsoft.com/default.asp.

- **Programming Collaborative Web Applications with Microsoft Exchange 2000 Server**

  Teaches Visual Basic programmers how to build applications using the Web Storage System and Exchange 2000 Server. By Mindy Martin for Microsoft Press, available Summer 2000. Visit http:l/mspress.microsoft.com/developer/.

- **Programming Microsoft Outlook and Exchange 2000, Second Edition**

  Describes how to develop applications for Outlook and Exchange using new Outlook 2000 features, such as Team Folders and Form Printing. By Thomas Rizzo for Microsoft Press, available Summer 2000. Visit http://mspress.rnicrosoft.com/developer/.

Enterprise Deployment Guide Contents

PART 1 Exchange 2000 Project Planning 1

CHAPTER1 What's New 3

Flexible Administration Model 3

Active Directory 4

Client Directory Access 4

Design Considerations 5

Administration 5

Administrative Groups 6

Permissions 6

MMC 6

Policies 6

Message Routing 7

Routing Groups 7

SMTP Transport 7

Link State Algorithm 8

Routing Group Connectors 8

SMTP Connector 9

Connectivity to Other E–Mail Systems 9

Enhanced Data Storage 10

Multiple Databases Per Server 10

Web Support 10

Complete Client and Protocol Support 11

Integrated Full–Text Search 11

CHAPTER 1 What's New *(continued)*

Scalability and Reliability . . . 12
Real–Time Collaboration . . . 12
Chat Service . . . 12
Instant Messaging Service . . . 13
Exchange 2000 Conferencing Server . . . 13
Conference Management Service . . . 13
Data Conferencing Provider . . . 13
Video Conferencing Provider . . . 14
Third–Partysupport . . . 14
Custom Solution Development . . . 14
XML Support . . . 14
Enhanced CDO . . . 14
Server Events . . . 15
Enhanced Workflow . . . 15
OLE DB and ADO Support . . . 16
IIS and ASP Integration . . . 16
Web Forms . . . 17

CHAPTER 2 Buildingthe Project Plan . . . 19

Taking a Phased Approach . . . 20
Assessing Project Risk . . . 20
Example . . . 21
Templates to Aid with Your Project . . . 22
Phase I: Envisioning . . . 22
Building the Project Team . . . 23
Building the Project Structure . . . 25
Gathering High-Level Requirements . . . 25
Defining the Project Vision . . . 25
Defining the Project Scope . . . 26
Defining the Project Assumptions . . . 26
Creating a Conceptual Design . . . 27

CHAPTER 2 Building the Project Plan *(continued)*

Phase II: Planning 27
Gathering Information 28
The Functional Specification 29
Proof-of-Concept Testing 29
Identifying Resources 30
Creating Required Project Plans 30
Building the Project Schedule 31
Phase III: Developing 32
Validating the Designs 32
Building the System 33
Pre-Pilot Tests 33
User Pilot Test 33
Phase IV: Deploying 34
Migrating Users 34
Handing Off the Project to Production 35
Completing the Project 35
PART 2 Planning for Exchange 2000 and Active Directory 37
CHAPTER 3 The Exchange 2000 Environment 39
Company Structure 40
Business Requirements 41
Organizational Requirements 41
User Requirements 42
Administrative Requirements 42
Future Planning 43
Current Network Infrastructure 43
WAN Environment 44
LAN Environment 45
Domain Name System 46
Windows NT 4.0 Domain Structure 46

CHAPTER 3 The Exchange 2000 Environment *(continued)*

Active Directory 47
Forests 47
Sites 47
Schema 48
Global Catalog and Global Catalog Servers 48
Active Directory Connector 49
Exchange Architectural Review 49
Overall Exchange Design 50
Address Book Views 51
Exchange Site Design 51
Mailbox and Public Folder Servers 51
Bridgehead Servers 52
Connector Servers 53
Backbone 53
Internet Mail 53
Disaster Recovery Planning 54
Exchange Monitoring and System Performance 55
Server Configuration 55
Hardware 55
Clustering 56
Windows2000 57
Exchange Software Version 57
Third-Party Software 57
Server Names 58
Client Access 58
Public Folders 59
Public Folder Hierarchy 59
Public Folder Affinity 59
Replication 60
Full Content Indexing 61

CHAPTER 3 The Exchange 2000 Environment *(continued)*

Organizational Forms 61
Conclusion 62
CHAPTER 4 Active Directory Design 63
Active Directory Overview 63
Namespaces in Your Company 64
Planning the DNS Namespace 65
Defining the Namespace Architecture 66
Domain Structure 66
Domain Controller 66
Domain Mode 68
Domain Tree 69
Domain Name System Overview 70
Resource Records 70
Zones 71
Active Directory Logical Structure 71
Domain Design 72
Determining the Number of Domains in Each Forest 72
Choosing a Forest Root Domain 74
Changing the Domain Plan After Deployment 75
Exchange 2000 and Active Directory Domain Design 75
DNS Service 76
Advantages of DNS Service 76
Using Non-Windows 2000 DNS Service 77
Placing DNS Servers 78
Domain Naming Recommendations 79
DNS and Exchange 2000 81
User Principal Name 83

CHAPTER 4 Active Directory Design *(continued)*

Organizational Unit Structure 84
Organizational Unit Characteristics 84
Organizational Unit Planning Process 85
Exchange 2000 and Organizational Units 87
Forest 88
Number of Forests 88
Exchange 2000 and the Forest 89
Windows 2000 Site Topology 90
Distinguishing Windows 2000 Sites from Exchange 5.5 Sites 90
Active Directory Logical and Physical Structure 91
Creating a Windows 2000 Site Topology 91
Site Links 92
Exchange 2000 and Windows 2000 Site Design 93
Trust Relationships 94
Transitive and Non-transitive Trust 94
Limiting Trust Relationships 95
Optimizing Authentication with Shortcut Trust Relationships 96
Active Directory Logical Components 96
Naming Contexts 96
Domain Naming Context 96
Configuration Naming Context 97
Schema Naming Context 97
Exchange 2000 and Naming Contexts 97
Global Catalog Server 98
Global Catalog Server Placement 98
Exchange 2000 and Global Catalog Server 99
Global Address List 99

CHAPTER 4 Active Directory Design *(continued)*

Active Directory Groups 100
Group Types 100
Group Scopes 101
Mail-Enabling Groups 103
Exchange 2000 and Groups 104

CHAPTER 5 Active Directory Integration and Replication 109

Active Directory Replication 110
Global Catalog Servers 111
Determining the Number of Global Catalog Servers Required 112
Global Catalog Database Sizing 112
Global Catalog Server Sizing 114
Global Catalog Server Placement 114
Selectable Field Replication 115
Selecting Attributes to Replicate to the Global Catalog 117
Exchange 2000 Installation and Active Directory Schema 117
Making Schema Changes 119
Address Book Searches 119
Active Directory Services 120
Global Address List 120
Exchange Server 5.5 Distribution Lists and Active Directory Groups 120
Address Lists and Offline Address Lists 121
Address Book Views 121
Address Lists 121
Address List Compatibility with Exchange Server 5.x 122
Recipient Update Service 123
Offline Address Lists 123
Client Access to Active Directory 124
DSProxy Service 124
Exchange Client. Outlook 97, and Outlook 98 126
Outlook 2000 and Outlook 98 SP2 126

CHAPTER 5 Active Directory Integration and Replication *(continued)*

Coexistence and Upgrading 127
What is the ADC? 127
Site Replication Services and Recipient Update Service 128
Why Use the ADC? 129
Connection Agreements 130
User Connection Agreements 131
Configuration Connection Agreements 133
Creating Initial Connection Agreements 134
ADC Operation 134
ADC Object Matching 135
ADC Replication 137
ADC Installation Location 138
Questions to Ask Before ADC Deployment 139
Preparing to Deploy the ADC 140
Exchange 5.5 Directory Preparation 140
Server and System Preparation 141
Accounts Preparation 141
Contingency Plan Preparation 142
Exchange Organization Preparation 142
Minimum Requirements for Installing the ADC 145
Monitoring the ADC 146
Deployment Tips 148

CHAPTER 6 Deployment Strategies 149

Deployment Roadmap 149
New Exchange 2000 Organization 150
Existing Exchange Organization 150
Preparing to Deploy 151
Deciding Whether to Install ADC First or Run ForestPrep 152

CHAPTER 6 Deployment Strategies *(continued)*

Populating User Accounts in Active Directory . . . 153
Scenarios . . . 154
In-Place Domain Upgrade Followed by ADC . . . 154
Run ADC, Upgrade Later, and Run Active Directory Cleanup Wizard . . . 154
Run ADC and Migrate Later By Using the Active Directory Migration Tool . . . 155
Run Active Directory Migration Tool with sIDHistory and Then Run ADC . . . 156
Run Active Directory Migration Tool Without sIDHistory and Then Run ADC . . . 157
Upgrading Windows NT Server to Windows 2000 Server . . . 157
Do You Need to Upgrade All Servers? . . . 158
Upgrade an Existing Domain and Connect to a Forest . . . 158
Design Issues . . . 159
Exchange Running on Primary or Backup Domain Controllers . . . 159
Upgrade Servers and Migrate to a Forest . . . 160
Upgrading Exchange Mailboxes . . . 161
Coexistence . . . 163
Configuration Connection Agreements and Site Replication Service . . . 164
Distribution Lists and Permissions . . . 170
Services That Cannot Be Upgraded . . . 171
Scenarios . . . 171
Separate Exchange Organizations . . . 171
Exchange 2000 in Its Own Forest . . . 172
Multiple Exchange 5.5 Organizations . . . 174

PART 3 Prototyping Exchange 2000 . . . 175

CHAPTER 7 Setting Up a Test Environment . . . 177

Why Testing is Important . . . 177
Planning the Test Lab . . . 178
Office Space . . . 179
Hardware . . . 179
Networking . . . 180
Software . . . 180

CHAPTER 7 Setting Up a Test Environment *(continued)*

Gathering Requirements . . . 181
Developing a Test Plan . . . 181
Identify Risks and Contingencies . . . 181
Network Effects . . . 181
Administrative Access . . . 182
SMTP Connector . . . 182
Replication Latency . . . 183
Training . . . 183
Third-Partyconnectors . . . 183
Client Software Test Plan . . . 183
Third-Party Connectors Test Plan . . . 185
Test Strategy . . . 186
Testing Your Scenarios . . . 186
Feature Tests . . . 186
Multiple Databases . . . 187
Storage Groups . . . 187
Backup and Recovery . . . 187
Public Folders . . . 188
Full-Text Indexing . . . 188
Address Lists . . . 188
Routing Groups . . . 188
Message Transport Between Routing Groups . . . 189
Administrative Groups . . . 189
Policies . . . 189
Distribution List Management and Usage . . . 190
Web Folders . . . 190
Upgrade Testing . . . 190

CHAPTER 8 Piloting Exchange 2000 . . . . . **193**
The Role of the Pilot 194
Determining Pilot Objectives 195
Choosing Which Features to Pilot 196
Existing Infrastructure Requirements 197
Server-Side Features 197
Client-Side Features 197
Who Should Participate? 198
Setting User Expectations 198
How Long Will a Pilot Take? 199
Production Pilots 199
Moving From Lab to Pilot 199
Documenting Lab Configurations 199
Evaluating Lessons from Lab Tests 200
Documenting Pilot Processes 200
Applying Lessons to Production 200
CHAPTER 9 Preparing a New Environment **201**
Assess and Evaluate Environment 202
Define the Required Functionality 202
Design Goals 203
Conduct a Gap Analysis 203
Create the Plan 205
Prepare the Environment 207
Physical Network 207
Windows 2000 Global Catalog Considerations 208
End-User Input 209
Design Exchange 2000 Architecture 210
Server Locations 210
Connectivity from the Retail Stores and the Internet 213

CHAPTER 9 Preparing a New Environment *(continued)*

Global Catalog Placement 214
Message Routing 214
Message Routing Option 1 215
Message Routing Option 2 216
Message Routing Option 3 217
Connectivity to the Internet 219
Real-Time Collaboration Services 219
Data Conferencing Provider 220
Server Roles 221
Mailbox Servers 221
Public Folders Servers 223
Connector Servers 224
Front-Endservers 224
Data Conferencing Servers 225
Naming Conventions 225
Exchange Organization Naming 226
Routing Group Naming 226
Administrative Group Naming 226
Server Naming 227
SMTP Alias Format 227
Mail-Enabled Contacts for Non–LitWare Users 227
Mail-Enabled Contacts for LitWare Users 228
Lotus Notes Addresses 228
Mail-Enabled Groups 228
Mandatory Distribution Lists 229
Scheduling Resources As Recipients 229
Virtual Resources 230

CHAPTER 9 Preparing a New Environment *[continued]*

Backup and Restore 230
Administration 230
Public Folders 231
Public Folder Affinity 231
Future Public Folder Considerations 233
Client Access 233
Outlook2000 234
Outlook Web Access 234
POP3 and IMAP4 234
NetMeeting 234
Design Coexistence and Migration Plan 234
Coexistence 235
Coexistence Architecture 236
Migration 239
Special Considerations 239
Coexistence and Migration Process 240
Connecting the Gateways 240
Configuring the Internet Backbone 240
Migrating 240
Train Administrators and Users 241
Administrator Training 241
End-User Training 241
Design Summary 242

CHAPTER 10 Preparing an Existing Environment 245

Operating System Dependencies 246

Exchange 2000 Requirements 246

Hardware Requirements 247

Windows 2000 Server Infrastructure 247

Compatible Versions 248

Active Directory 248

DNS Service 251

NNTP Service 251

SMTP Service 251

Optional Windows 2000 Services 251

Exchange 4.0. Exchange 5.0. and Exchange 5.5 Dependencies 251

Upgrade Paths 252

Migration Options 252

Installation Considerations 252

Disk Space 252

Time Required to Upgrade a Server 252

Permissions Required to Upgrade 253

Windows Trusts 253

Server Roles 253

Mixed and Native Modes 253

Routing and Routing Masters 254

Offline Address Books 255

Schedule+ Free/Busy and Organizational Forms Library Public Folders 255

Exchange Component Comparison 255

CHAPTER 10 Preparing an Existing Environment *(continued)*

Directory Objects 257
Mailbox–Enabled and Mail–Enabled Objects 257
Custom Recipients 257
Distribution Lists 257
Recipient Containers 257
Address Book View Containers 258
Address Book Views 258
Hidden Objects in the Exchange 5.5 Directory 258
Names of Exchange 5.5 Custom Attributes 258
Web Storage System 259
Mailboxes 259
Move Mailbox Function 259
Public Folders 259
Circular Logging 262
Communication Protocols 262
Protocol Servers 262
Outlook Web Access 263
NNTP 263
IMAP4 263
POP3 263
LDAP 264
Link Monitors 264
Server Monitors 264
Message Tracking 264
Tracking Log Format 265
Message Event Log Format 265

CHAPTER 10 Preparing an Existing Environment *(continued)*

Transport Stacks and Connectors 265
Remote Access Service 265
TP4 265
X.400 Connectors 266
Site Connectors 269
Dynamic RAS Connector 270
Internet Mail Connector 270
PROFS Connector 275
SNADS Connector 275
Other Exchange Services 275
Chat Services 275
Key Management Service 276
Event Scripts 276
Third–Party Software 276
Client Effects 276
Exchange Client. Outlook 97 and Outlook 98 276
Outlook2000 277
Outlook Profiles 277
POP3 and IMAP4 Clients 277

PART 4 Basic Deployment Planning 279

CHAPTER 11 Administration and Maintenance 281

Recipient Management 282
Organizational Units 282
Administrative Models 283
Centralized Administration 283
Delegated Administration 284
Messaging Group Administration of Exchange Attributes 284

CHAPTER 11 Administration and Maintenance *(continued)*

Understanding Administrative Tools and Features for Recipient Management 284
Administrative Tools 284
Organizational Units 285
Active Directory Delegation of Control Wizard 285
Exchange Administration Delegation Wizard 286
Permissions 286
Recipient Policies 287
Server Management 287
Managing Server Responsibilities 287
Administrative Models 288
Centralized Management 289
Distributed Management 289
Mixed Management 290
Understanding Administrative Tools and Features for Server Management 290
Domains 290
Organizational Units 290
Administrative Groups 290
Routing Groups 291
System Policies 291
Exchange Administration Delegation Wizard 292
Additional Administrative Features 292
Security and Distribution Groups 293
Public Folders 293
Access Control 293
Public Folder Affinity 294
Moving Mailboxes 294
Additional Hardware 295
Network Bandwidth 295
Group Policy 296

CHAPTER 11 Administration and Maintenance *(continued)*

Mixed Mode Operation with Exchange 5.5 . . . 297
Administrative Groups . . . 297
Active Directory Connector . . . 298
Public Folders . . . 298
Moving an Exchange 5.5 Server . . . 298
Managing and Monitoring Exchange Server Performance . . . 299
Exchange 2000 Tools . . . 300
Exchange Monitoring and Status Tool . . . 300
Queue Viewer . . . 302
Message Tracking Center . . . 302
Windows 2000 Tools . . . 302
Performance Monitoring . . . 303
Event Viewer . . . 303

CHAPTER 12 Sewer Design for Backup and Restore . . . 305

Databases and Storage Groups . . . 305
Backup . . . 307
Restore . . . 307
Parallel Operations . . . 308
Server Recovery . . . 309
Recommendations . . . 309
Design Issues . . . 309
Backup to Disk . . . 309
Backup and Restore Throughput . . . 310
Sample Designs . . . 311
Small Organization or Branch Office Servers . . . 311
Large Departmental Server . . . 312
Centralized Server or Data Center Installation . . . 315

CHAPTER 13 Virus Protection 319
Virus Protection Overview 319
Virus Scanning Concepts 321
Definitions 322
Solutions Planning 322
Ties 1: Backbone Infrastructure 323
Tier 2: Local Servers 324
Tier 3: Desktop 324
Virus Scanning Support in Exchange 325
Virus Scanning API in Exchange 5.5 SP3 325
Installation 325
Operation 326
Administration 326
Performance 326
Logging Features 327
Virus Scanning Support in Exchange 2000 Server 327
Protection Without the Virus Scanning API 327
High Performance with the Virus Scanning API 327
Encrypted Messages 327
Virus Scanning Support in Outlook 328
Spam Support in Outlook 328
Virus Solution Vendors 329
Detection Circumvention Tests 330
Performance Objects 331
Monitoring Service 332
Alert Monitoring 332

CHAPTER 14 Server Availability **333**
Defining High Availability Architecture 333
Determining Availability Requirements 334
Cost of Downtime 335
System Vulnerability 336
Recovery Point and Recovery Time 337
Causes of Downtime 338
Implementing Availability Technologies 339
Server Hardware 339
Distributed Services 339
Multiple Databases 340
Intelligent Routing 342
Server Clusters 344
Front-End and Back-End Architecture 345
Creating Availability Processes 348
Conclusion 349
CHAPTER 15 Server Sizing **351**
Exchange 2000 Server Types 352
Processor Configuration 353
Capacity for Mailbox and Public Folder Servers 354
Capacity for Connector Servers 354
Capacity for Data and Video Conferencing Servers 355
Capacity for Instant Messaging Servers 355
Capacity for Chat Servers 355

CHAPTER 15 Server Sizing *(continued)*

Disk Configuration 356
Disk Configuration for Mailbox and Public Folder Servers 357
Location for Mailbox and Public Folder Stores 358
RAID for Mailbox and Public Folder Servers 359
Disk Controller Cache Settings for Mailbox and Public Folder Servers 359
Paging File Settings for Mailbox and Public Folder Servers 360
MIME Content 361
Index Size 361
Disk Space Margin of Safety 361
Site Replication Service 361
Disk Configuration for Connector Servers 362
Location for Connector Servers 366
RAID for Connector Servers 367
Disk Controller Cache Settings for Connector Servers 367
Antivirus Solutions 367
Disk Configuration for Front-End Servers 367
Disk Configuration for Data and Videoconferencing Servers 368
Disk Configuration for Instant Messaging Servers 368
Disk Configuration for Chat Servers 369
Memory Configuration 370
Memory for Mailbox and Public Folder Servers 370
Memory for Connector Servers 370
Memory for Front-End Servers 371
Memory for Data and Video Conferencing Servers 371
Memory for Instant Messaging Servers 371
Memory for Chat Servers 372
Network Configuration 372
Network Configuration for Mailbox and Public Folder Servers 372
Network Configuration for Connector Servers 372
Network Configuration for Real-Time Collaboration Servers 372

CHAPTER16 Message Routing **373**
Exchange 5.5 Routing Basics 373
Exchange 2000 Routing Basics 374
Administration Groups and Routing Groups 375
Connectors 375
Routing Group Connector 375
SMTP Connector 376
X.400 Connector 376
Link State Routing 377
Route Selection 379
Coexistence 379
Topology Considerations 380
CHAPTER 17 Backbone Configuration and Tuning **383**
Directory Access 384
Directory Service Access API 384
Specifying an Active Directory Server 384
Directory Cache 386
Cache Parameters 386
Routing and Transport 387
Advanced Queuing Engine 387
Message Categorizer 387
Distribution Group Expansion 388
Connector Types 388
Routing Group Connector 389
SMTP Connector 390
X.400 Connector 390
Routing and Link State Information 391
How Link State Works 391
Link State Updates 392
Message Delivery to Remote Servers 393

CHAPTER 17 Backbone Configuration and Tuning *(continued)*

Routing Group Scenario 393
Low-Bandwidth Environments 395
Connector Performance 395
Non-Connected Networks 396
Domain Controller and Global Catalog Server Placement 396
Public Folders 396
Public Folder Hierarchies 397
Configuring Public Folders 398
Public Folders in Active Directory 398
Public Folder Replication 399
Public Folder Affinity 399
Address Lists 400
Address List Compatibility with Earlier Versions 401
Recipient Update Service 401
CHAPTER 18 External Connectivity 403
Planning and Best Practices 403
Active Directory and Connectivity 407
Objects 407
Global Catalog 408
Recipient Policies 408
Recipient Update Service 409
Connector Review 409
Exchange Lotus Notes Connector 409
Messaging 410
Directory Synchronization 410
Configuration 411

CHAPTER 18 External Connectivity *(continued)*

Exchange Microsoft Mail Connector . . . 414
Messaging . . . 414
Directory Synchronization . . . 414
Calendaring . . . 416
Configuration . . . 416
Exchange Lotus cc:Mail Connector . . . 417
Scenario . . . 419
Configuration . . . 420
Exchange Connector for Novell GroupWise . . . 422
Messaging . . . 422
Directory Synchronization . . . 422
Configuration . . . 423
Connectivity in Mixed Mode . . . 424
Active Directory Connector . . . 425
Routing in Mixed Environment . . . 425
Connector Migration . . . 425
Connectivity to PROFS and SNADS . . . 426
Directory Synchronization with SNADS and PROFS Connectors . . . 427
Microsoft Metadirectory Services . . . 427

PART 5 Advanced Deployment Planning . . . 429

CHAPTER 19 Chat and Instant Messaging Services . . . 431

Chat Service Overview . . . 431
Deploying Chat Service . . . 432
Chat Service Deployment Process . . . 433
Chat Server Scalability . . . 434
Chat Server Location . . . 434
Network Address Translation and Internal Firewalls . . . 436
Security and TCP Ports . . . 436
Basic Chat Service Recommendations . . . 437
Instant Messaging Overview . . . 437

CHAPTER 19 Chat and Instant Messaging Services *(continued)*

Deploying Instant Messaging 438
Instant Messaging Deployment Process 439
Exchange 2000-Only Deployment 439
Exchange 5.5 Deployment 440
Instant Messaging Routers and Home Servers 441
Relaying Requests Across the Network 444
Resolving Names for Instant Messaging Servers 445
Registering Instant Messaging Routers in DNS 445
Client Name Resolution 445
Scalability for Instant Messaging Servers 446
Server Location and Network Connectivity 446
Low Bandwidth Connections 447
Branch Offices 447
Network Address Translation and Internal Firewalls 448
Security Considerations 448
Instant Messaging Deployment Recommendations 448

CHAPTER 20 Inter-Organization Replication and Directoty Synchronization 449

Scenarios 450
Exchange 5.5 Inter-Organization Solutions 450
InterOrg Synchronization Tool 451
Windows 2000 with Active Directory Connector 451
Configuring ADC for Inter-Organization Synchronization 452
Versions of ADC 453
Exchange 2000 Inter-Organization Solutions 453
Microsoft Metadirectory Services 454
Active Directory Management Agent 457
Inter-Forest Scenarios 458
Centrally Managed Forests 458
Peer Forests 459
Conclusion 461

CHAPTER21 Branch Office Scenarios 463

Replication and Routing Group Dependencies in Windows 2000 463

Exchange 2000 Routing Groups and Administrative Groups 464

Designing Routing Groups 465

Branch Office Administrative Models 467

Planning Your Exchange 2000 Administrative Model 467

Present Administration Model 468

Network Considerations 468

Different Types of Mail Client Access 468

Security 468

Active Directory Permissions 469

Centralized Exchange 2000 Administrative Model 469

Routing Group Strategies 469

Administrative Group Strategies 472

Windows 2000 Account Administration and Exchange 2000 Mailbox Administration 472

Distributed Exchange 2000 Administrative Model 473

Routing Group Strategies 473

Administrative Group Strategies 474

Windows 2000 Account Administration and Exchange 2000 Mailbox Administration Strategies 475

Mixed Centralized and Distributed Exchange 2000 Administrative Model 475

Routing Group Strategies 475

Administrative Groups 477

Windows 2000 Account Administration and Exchange 2000 Mailbox Administration 477

Global Catalog Placement 477

CHAPTER 21 Branch Office Scenarios *(continued)*

Branch Office Messaging Client Considerations 478
Deployment Objectives 479
Functionality Requirements 479
Mobile vs. Offline Requirements 479
Outlook2000 479
IMAP4 and POP3 481
Outlook Web Access 483
Terminal Services 484
CHAPTER 22 Corporate Backbone Scenario 487
Financial Bank, Inc. 487
Using Exchange 2000 as a Backbone 488
Advanced SMTP Command Verbs 488
Chunking 488
Pipelining 489
The Web Storage System 489
Storage in Native Internet Format 489
Public Folders 490
Message Routing 490
Link State Algorithm 491
Advanced Queuing Engine 491
Message Categorizer 491
Exchange Instant Messaging Service 492
Chat Services 492
Data Conferencing Provider 492
Auditing Your System 493
Directory Synchronization 494
Message Routing and General Configuration 494
Internet Mail Architecture 494

CHAPTER 22 Corporate Backbone Scenario *(continued)*

Building the Directory Topology 496
Planning Active Directory 496
Active Directory Connector 496
Global Catalog Servers 498
Name Resolution 499
Directory Synchronization Between Forests 499
Building the Messaging Backbone Topology 500
Routing Groups 500
Link State Algorithm Considerations 501
Routing Group Deployment Scenarios 501
Message Routing and Group Expansion 503
Messaging Coexistence 503
Connectors to Other Mail Systems 503
Summary 504
CHAPTER 23 Hosted Service Environments 505
Hosting Architecture 505
Service Offering 506
Customer Division 506
Partition the Data 507
Windows 2000 and Active Directory 507
Recommended Active Directory Partitioning 508
Host Services 509
What Is the Target Company Size? 509
Mailbox Storage 509
Example Service Offerings 510
Naming Conventions 511
Global Address List 513
Public Folders 513

CHAPTER 23 Hosted Service Environments *(continued)*

Provide Shared Services 515
Provide Desktop Services 516
Administration 516
Configuration Checklist 516
Conclusion 517
CHAPTER 24 Security Sensitive Environments 519
Digital Encryption and Signatures 520
E-mail Security 521
Message Body Encryption 521
Digital Signatures 523
Certificates and Certification Authorities 523
Certification Authority 524
Certificate Enrollment 524
Certificate Revocation 525
Key Management Service 525
Cryptographic Services 526
Key Management Service Passwords 526
Key Management Service Process Flow 526
Enrollment 527
Certificate Renewal 530
Key Recovery 531
Certificate Revocation 531
S/MIME Design Scenarios 532
Intra-Company Scenario 532
Considerations for a Commercial Root CA vs. Self-hosted Root CA 534
Building a Public Key Infrastructure 534
Inter-Companyscenario 537
Building a PKI Between Companies 538
Sharing Directory Information 538

CHAPTER 24 Security Sensitive Environments *(continued)*

Inter-Company with Trusted Third Party scenario . . . 541
Building a PKI With a Third Party . . . 541
Deployment Scenarios . . . 546
Intra-Company E-mail Security . . . 546
Intra–Company Email Security Scenario . . . 547
Intra-Company Email Security Scenario . . . 548
Intra-Company E-mail Security Scenario . . . 549
Inter-Company E-mail Security . . . 550
Directory Replication Between Two Forests . . . 551
Smart Cards . . . 552
Trusted Third Party (Internet Directory Services) . . . 552
Replicating Directory Information to a Trusted Third Party . . . 554
Accessing Directory Information in a Trusted Third Party . . . 554
Building Certificate Trust Lists for Each Company . . . 554
Publishing Certificate Revocation Lists to a Trusted Third Party . . . 554

CHAPTER25 Outlook Web Access . . . 557

Outlook Web Access Evolution . . . 558
Features . . . 559
Accessing Your Mailbox . . . 559
Accessing Exchange Objects . . . 560
WebDAV and XML . . . 560
Clients . . . 561
Internet Explorer 5.0 . . . 561
Other Clients . . . 561
Client Limitations . . . 562
Deployment Planning . . . 562
Securing Outlook Web Access on the Internet . . . 563
Load Balancing and Fault Tolerance . . . 564

CHAPTER 25 Outlook Web Access *(continued)*

Front-End and Back-End Server Architecture 565
Configuring Back-End Servers 565
Configuring Front-End Servers 565
Configuring Network Load Balancing 566
Installing Exchange 2000 Server 567
Configuring DNS Entries 567
Authentication 567
Basic Authentication 568
Integrated Windows Authentication 568
Secure Sockets Layer 568
Anonymous Access 568
Front-End and Back-End Authentication 569
Pass-Through Authentication 569
Dual Authentication 569
Performance Monitoring 569
Applications 570
APPENDIX A Client Network Traffic Analysis 573
Test Lab Configuration 574
LAN Design 576
Servers Characteristics 576
Client Characteristics 577
Exchange 2000 Configuration 578
Terminal Services Configuration 579
Client Configuration 579
Outlook 2000 Client Configuration 579
Outlook 97 Client Configuration 579
Outlook Express Client Configuration 580
Netscape Messenger Client Configuration 581
Outlook Web Access Client Configuration 582
Terminal Services Client Configuration 582

APPENDIX A Client Network Traffic Analysis *(continued)*

Measurement Methodology . . . 583
Log On and Log Off . . . 584
Tests Performed . . . 585
MAPI Clients: Microsoft Outlook 2000. Outlook 97 . . . 585
Test Details . . . 585
Log On and Log Off Results . . . 588
Outlook 2000 and Outlook 97 Measurements Analysis . . . 588
Microsoft Outlook Express: POP3 and IMAP4 Modes . . . 588
Testdetails . . . 588
Outlook Express IMAP4 and POP3 Results . . . 590
Outlook Express Measurements Analysis . . . 590
Netscape Messenger: POP3 & IMAP4 Modes . . . 590
Test Details . . . 591
Netscape Messenger IMAP4 Results . . . 591
Netscape Messenger IMAP4 Measurements Analysis . . . 591
Outlook Web Access . . . 591
Test Details . . . 591
Outlook Web Access Results . . . 592
Outlook Web Access Measurements Analysis . . . 592
Directory Access . . . 592
Tests Performed . . . 593
MAPI Clients: Outlook 2000. Outlook 97 . . . 594
Test Details . . . 594
Outlook 2000 and Outlook 97 Results . . . 599
Outlook 2000 and Outlook 97 Measurements Analysis . . . 599

APPENDIX A Client Network Traffic Analysis *(continued)*

Outlook Express: LDAP mode 600
Test Details 600
Outlook Express LDAP Results 601
Netscape Messenger: LDAP mode 601
Outlook Web Access 601
Test Details 601
Outlook Web Access Results 602
Outlook Web Access Measurements Analysis 602
Mail Items 603
Tests Performed 603
Test Details 603
Outlook 2000 Results 608
Outlook 97 Results 609
Outlook Express IMAP Results 609
Outlook Express POP Results 609
Netscape Messenger IMAP Results 610
Netscape Messenger POP Results 610
Outlook Web Access 611
Test Details 611
Outlook Web Access Results 614
Mail Item Analysis 614
Calendaring. Contacts and Tasks 615
Tests Performed 615
Test Details 616
Outlook 2000. Outlook 97. and Outlook Web Access Results 618
Calendaring. Contacts. and Tasks Analysis 618

APPENDIX A Client Network Traffic Analysis *(continued)*

Public Folders . . . 619
Tests Performed . . . 619
Test Details . . . 620
Outlook 2000 Results . . . 621
Outlook 97 Results . . . 621
Outlook Express NNTP Results . . . 621
Netscape Messenger NNTP Results . . . 622
Outlook Web Access Results . . . 622
Public Folder Analysis . . . 622
Outlook 2000 with Terminal Services . . . 623
Test Details . . . 623
Connect . . . 623
TS Logon . . . 624
TS Logoff . . . 624
Outlook Logon . . . 624
Outlook Logoff . . . 625
Message Actions . . . 625
Calendar Actions . . . 626
Contact Actions . . . 626
Task Actions . . . 627
Note Actions . . . 628
Address Book Actions . . . 628
Terminal Services Client Results . . . 629
Terminal Services Measurement Analysis . . . 629
Web Storage System . . . 630
Tests Performed . . . 630
Log On and Log Off and User Traffic Results . . . 631
Web Storage System Measurement Analysis . . . 631

APPENDIX A Client Network Traffic Analysis *(continued)*

Instant Messaging 631
Tests Performed 631
Log On and Log Off and User Traffic Results 632
Status Change 632
Sending Messages 632
Client Traffic Measurement Conclusions 633
DSProxy 634
Outlook 98 and Outlook 2000 635
Exchange Client and Outlook 97 and Outlook 98 636
Traffic and Load Generated Through the DSProxy Process 637

PART 1

Exchange 2000 Project Planning

In This Part

What's New

Building the Project Plan



CHAPTER 1

What's New

Robert Dring, Senior Consultant, Microsoft

Microsoft Exchange 2000 Server includes features for large scale messaging, collaboration, and custom solution development. Some of these features are new to the product; some are improvements on features in Microsoft Exchange Server version *5.5;* and some are major changes to old features. Despite these changes, all clients that work with Exchange Server *5.5* and earlier also work with Exchange 2000 Server—without any modifications.

The following discussion introduces the major features of Exchange 2000 Server and outlines how they can enhance administration, connectivity, collaboration, and application development. Details about each of the topics in this chapter appear in later chapters.

In This Chapter

Flexible Administration Model

Connectivity to Other E–Mail Systems

Enhanced Data Storage

Scalability and Reliability

Real–Time Collaboration

Custom Solution Development

Flexible Administration Model

In an installation of Exchange *5.5,* servers are grouped into one or more sites. An Exchange site is a group of servers that are administered as a unit, that route mail directly to each other, and that replicate directory information from one server directly to another. The administrator must configure explicit message routing and directory replication connections between Exchange sites. However, in an Exchange *5.5* organization, two opposing needs can conflict: on one hand, you may want to have large sites for administration purposes; on the other hand, the network may not have sufficient bandwidth for message routing and directory replication throughout a large site.

In Exchange 2000, the following three functions of an Exchange 5.5 site are separated:

- **Directory replication** Directory replication is now handled by the Microsoft Windows 2000 operating system, independently of Exchange. A deployment of Windows 2000 can include sites, which are a collection of IP subnets with high–bandwidth connections between them, but these sites are not the same as Exchange sites.
- **Administration** Exchange 2000 servers are organized into administrative groups for management purposes.
- **Message routing** Exchange 2000 servers are organized into routing groups to control message flow. There is no requirement for the routing group and administrative group organization to be the same.

With Exchange 2000, you can now place servers into administrative groups that match your operational structure. You can also deploy Windows 2000 domain controllers in Windows 2000 sites and domains to match IP subnets and Active Directory operational structure. Finally, you can leave directory management to the Windows 2000 administration team.

Note that while Exchange 2000 coexists with Exchange 5.5, administrative groups and routing groups include the same servers and show up in the Exchange 5.5 directory as Exchange sites. After all the Exchange 5.5 servers are upgraded to Excnange 2000 and the Exchange organization is changed to native mode, you can define routing groups that do not correspond one to one with administrative groups. Even if Exchange 2000 does not coexist with earlier versions, it installs by default in mixed mode; you must switch it manually to native mode.

Active Directory

Exchange 2000 uses the Active Directory directory service in Windows 2000, which provides lower cost of ownership—one directory to maintain instead of two—and a unified security model. In Exchange 5.5, each server in the organization maintains its own copy of the Exchange directory and validates users' security permissions using the separate directory in Microsoft Windows NT Server version 4.0.

In Exchange 2000, permissions are defined using objects in Active Directory. Some permissions in Exchange 5.5, such as mailbox ownership and Exchange administration rights, are specified by using Windows NT accounts and groups; other permissions use Exchange mailboxes and distribution lists.

Client Directory Access

With Exchange 2000, clients access the directory on Windows 2000 global catalog servers. With Exchange Server 5.5 and earlier, clients access the directory on the mailbox server.

Microsoft Outlook 2000 and Outlook 98 Service Pack 2 can access the global catalog servers directly. Earlier clients still direct their requests to their mailbox server, which redirects them to a global catalog server. This redirection capability is a key part of the support that Exchange 2000 provides for earlier clients.

Design Considerations

Although Exchange 2000 has built-in support for earlier clients, there are still some issues that Exchange administrators need to be aware of. It is important that Exchange administrators be part of the Active Directory planning effort. Indeed, because Exchange 2000 is the first enterprise-wide application that requires Active Directory, it is likely that Exchange administrators will drive Active Directory deployment. The following design issues for Active Directory might be of interest to Exchange administrators.

Global catalog replication Exchange clients display directory information that is obtained from global catalog servers. The global catalog includes every object in the entire Windows 2000 forest but only selected attributes for each object replicates to global catalog servers. You must ensure that the attributes you rely on in the Exchange 5.5 directory are also present in the global catalog.

Windows 2000 groups Mail-enabled groups in Active Directory replace the distribution lists provided by Exchange 5.5. Although all group types work for the purpose of mail routing, only universal group membership is viewable in all domains. This fact affects whether users can expand group membership—something that they can always do in Exchange 5.5. Furthermore, whereas both security and distribution groups can be used for mail routing purposes, only security groups can be used to control access to resources such as Exchange 2000 public folders or Windows 2000 file shares. The universal security group is available only in a domain in native mode.

Different appearance of the directory As with Exchange 5.5, Exchange 2000 users see a local address list. If they investigate further, users see different subgroups of the directory, instead of the familiar structure in Exchange 5.5, which consists of organization, site, and recipient containers.

Delegation of user administration In Exchange 2000, parameters such as mailbox location and size limits are now just attributes of objects in Active Directory. The same person managing Windows 2000 accounts can manage these parameters.

Exchange 2000 allows you to reduce the cost of administration by integrating roles. Alternatively, you can maintain separate roles by specifically assigning permissions to manage a set of objects. For example, permissions on user objects or mailbox attributes can be assigned to administrators who deal only with mailboxes.

Delegation of server administration You can use resource domains with Windows 2000 and Exchange 2000. In addition, you can organize Exchange 2000 servers in Active Directory organizational units and delegate administration of those organizational units to the messaging administrators.

Administration

In Exchange 2000, administration does not dictate your messaging system design. Thanks to the flexibility provided by administrative groups, permissions, Microsoft Management Console (MMC), and policies, you can customize Exchange administration to your needs.

Administrative Groups

Servers can be assigned to administrative groups without consideration for message routing. This gives administrators great flexibility in managing servers, regardless of a server's physical location. Note that this is possible only after the Exchange 2000 organization is in native mode. When Exchange 2000 and Exchange 5.5 operate in mixed mode, administrative groups match Exchange 5.5 sites, and servers must be members of routing groups within their administrative group.

You can move mailboxes between mailbox stores in an Exchange organization. However, mailbox moves are restricted to mailbox stores within the same administrative group when operating in mixed mode. It is not possible to move a server to another administrative group.

Permissions

All Exchange 2000 configuration information is held in Active Directory. Because Active Directory supports attribute–level security permissions, administrator? have precise control over Exchange security configuration.

You can use the Exchange Administration Delegation Wizard to quickly apply new security configurations and delegate rights to particular administrative groups or to the entire organization.

MMC

MMC provides consoles for Exchange 2000 that replace the Exchange Administrator program used in earlier versions. Administrators can include only the consoles relevant to their particular tasks. System Manager is the console for managing all aspects of an Exchange organization.

In addition, Exchange administrators can customize MMC consoles so that other administrators have control over only necessary features.

Policies

You can use System Manager to apply policies to a collection of objects of the same class. This way, administrators can control and easily modify the configuration of groups of servers. For example, you can create policies that specify settings for mailbox size limits and deleted item retention. You can apply one policy to most servers and selectively apply a different policy to other servers. For example, you can designate special policies for servers accessed by top executives. Furthermore, policies can exist in different administrative groups from the servers to which the policies apply. This means that you can use policies to manage some server settings and delegate other aspects of server management.

Message Routing

Exchange 2000 introduces many changes in message routing:

- Servers are organized into routing groups instead of Exchange sites.
- Message transport now uses Simple Mail Transfer Protocol (SMTP) instead of remote procedure call (RPC).
- Servers use a link state algorithm to exchange current connector status.
- The Routing Group connector replaces the Site Connector.

Routing Groups

An Exchange 2000 routing group is a collection of servers that route mail directly to each other. There can be many routing groups within an organization, but each server can belong to only one. Routing groups exist within administrative groups. In Exchange 2000 native mode, it is possible to place a server in a routing group even if that server does not belong to the administrative group that the routing group belongs to. This permits delegation of control so that one set of administrators can manage message routing while another administrator manages other aspects of the servers.

Note that if all servers belong to a single group with full–time, high–bandwidth connectivity, no special routing group configuration is required. Unless routing groups are explicitly defined, the routing group does not appear in System Manager; instead, Connections appears under an administrative group.

If necessary, you can create and remove routing groups when the organization is in native mode. You can dramatically change the entire routing architecture for an organization without reinstalling Exchange.

SMTP Transport

Servers in an Exchange 2000 routing group transport messages from one server to another by using SMTP. The success of the Internet shows that SMTP is well suited to transferring messages across low–bandwidth links. Exchange 2000 Server takes advantage of performance enhancements available in service extension updates to the SMTP standard—RFC 1830 for transmission of large messages and binary MIME messages and RFC 2197 for including multiple SMTP commands in a single transmission.

Note that this reliance on SMTP means that Exchange 2000 Server requires the TCP/IP network protocol, as does Active Directory. Exchange clients may continue to use TCP/IP, Internetwork Packet Exchange (IPX) or NetBIOS Enhanced User Interface (NetBEUI).

SMTP and other protocols used by Exchange 2000 are provided by the Internet Information Services (IIS) component of Windows 2000 Server. In addition, Exchange uses the event sink architecture of Windows 2000 to add extensions to the base SMTP service provided by the operating system.

Although SMTP is used as the message transfer protocol between Exchange 2000 servers, it is not necessary to register Mail Exchanger (MX) records in Domain Name System (DNS) for all servers. Typically, MX records are required only for servers that will receive SMTP messages from computers outside the Exchange organization. Instead of using MX records for routing SMTP messages in an Exchange 2000 organization, the servers use the new link state algorithm.

Link State Algorithm

Exchange 2000 Server has a new routing algorithm, known as the link state algorithm. It is based on the Open Shortest Path First (OSPF) algorithm used by network routers. Within each routing group, one server is designated as a routing group master. When servers discover a routing change, they alert the routing group master, which propagates that change to all other servers in the routing group and to other routing groups. The result is that within a few minutes of a change, all servers in the organization have updated routing information.

With Exchange 5.5, servers have a static representation of routing options throughout the entire organization. One server in each Exchange site updates this view once every 24 hours. The only current information that servers have is for the connections they can make directly to their immediate neighbors. You may want to define multiple redundant routes between Exchange sites to ensure message transport in the event of failures, but servers are constrained because they do not have current information. For example, if both routes to a site are down, messages can bounce between the two servers that are connected to the remote site.

Routing Group Connectors

The Routing Group connector in Exchange 2000 replaces the Site Connector in Exchange 5.5; it is intended to be the principal connector between routing groups. The Routing Group connector uses SMTP and, because it does not route messages outside Exchange, does not require MX records in DNS. The Routing Group connector no longer identifies a single server as a bridgehead server; it can use one, several, or all servers in each routing group. This gives administrators greater control over connector design.

With Exchange 5.5, many administrators choose the X.400 connector over the Site Connector so that two servers in each site can be responsible for routing mail. This provides redundancy while limiting the servers that transport mail between sites. With Exchange 2000, a single Routing Group connector specifying multiple bridgehead servers replaces a pair of X.400 connectors.

SMTP Connector

The SMTP Connector in Exchange 2000 uses the SMTP service provided by IIS and supports the updated service extensions that are mentioned earlier in this discussion. It replaces the Internet Mail Service in previous versions of Exchange. If you use it as a connector between routing groups, it shares link state information with other Exchange 2000 servers; however, the Routing Group connector is usually a better choice.

Connectivity to Other E-Mail Systems

Other connectors in Exchange 2000 Server provide functionality similar to that available in Exchange 5.5 Service Pack 3.

X.400 Connector To improve interoperability with SMTP, the X.400 connector was updated in Exchange 5.5 SP3 to support Multipurpose Internet Mail Extensions (MIME) Internet X.400 Enhanced Relay (RFC 2156). This support was added to the message transfer agent (MTA) for Exchange 5.5 SP3 and is not supported in the Exchange 2000 MTA.

Microsoft Mail Connector The Microsoft Mail Connector supports bi-directional mail transfer, directory synchronization, and scheduling.

Lotus cc:Mail Connector The Lotus cc:Mail Connector supports bi-directional mail transfer, directory synchronization, and meeting requests.

Lotus Notes Connector The Lotus Notes Connector supports bi-directional mail transfer, directory synchronization, and scheduling. Exchange 2000 also includes the Lotus Notes Application Connector, which synchronizes Exchange public folders with Lotus Notes databases.

Novell GroupWise Connector The Novell GroupWise Connector supports bi-directional mail transfer, directory synchronization, and scheduling.

Microsoft Exchange Server 5.5 Enterprise Edition also includes connectors for OfficeVision. These connectors are not available in Exchange 2000 Server; however, if Exchange 2000 and Exchange 5.5 coexist in your organization, you can use the Exchange 5.5 connectors.

Enhanced Data Storage

The Microsoft Web Storage System supports multiple databases per server, ubiquitous client and protocol support, and integrated full–text search. These features make it possible to use Exchange 2000 as a document and data management platform.

Multiple Databases Per Server

In earlier versions of Exchange, servers have one database for e–mail (the private information store), and one database for shared folders (the public information store). Exchange 2000 now supports up to four storage groups on a single server, each holding up to five databases. This ability to host multiple databases provides several benefits:

- As with earlier versions of Exchange, databases can be backed up online. In Exchange 2000, each database can be backed up independently; however, because one set of transaction log files is maintained for each storage group, it makes sense to back up a storage group all at once. You can also restore a single database while the others remain online.
- With Exchange 5.5, the number of users on a server is limited primarily by the maximum database size that can be restored in an acceptable period of time. With Exchange 2000, server size can grow dramatically; what would have been one private information store in Exchange 5.5 can now be multiple mailbox stores.
- Storage groups can be created to serve different needs. For example, one storage group can contain information supplied by a Network News Transfer Protocol (NNTP) newsfeed while another storage group can contain mailbox stores. You can enable circular logging for the storage group that contains the newsfeed, and disable it for the storage group that contains mailbox stores.

Note that although one server can have multiple public folder stores, only one public folder store contains the All Public Folders tree that is accessible by Microsoft Outlook and other MAPI clients. You can use protocols supported by the Web Storage System to access the other public folder stores.

Web Support

Compared with earlier versions, Exchange 2000 provides vastly improved Web access to data. You can access every item in the Microsoft Web Storage System by using HTTP. Predictable and understandable URLs such as http:ll.servernamelusername/inbox and *http:llsewernamelusernamel* calendar make it easy to incorporate Exchange elements into Web pages and corporate portals. In addition, developers can easily incorporate Exchange features in applications by using Microsoft FrontPage 2000 and Microsoft Visual Studio.

Outlook Web Access for Exchange 2000 includes drag–and–drop functionality, rich–text editing, and new right–click menus; the end–user experience is similar to the Microsoft Outlook client. Although other browsers provide access to all Exchange data, only Internet Explorer 5 takes advantage of these advanced Outlook Web Access features.

Exchange 2000 supports an extension to HTTP 1.1 called Web Distributed Authoring and Versioning (WebDAV), which is defined in RFC 2518. WebDAV allows a client to manipulate a remote server as if it were a file system. In addition, Exchange 2000 makes extensive use of Extensible Markup Language (XML) to structure data provided to HTTP clients.

Complete Client and Protocol Support

Exchange 2000 introduces new streaming technology that allows content such as MIME messages, voice mail items, or video to be streamed in and out of the Web Storage System without intermediate conversions.

As in earlier versions of Exchange, clients can access data in the Web Storage System by using MAPI, Post Office Protocol version 3 (POP?), Internet Message Access Protocol version 4 rev1 (IMAP4) and NNTP. New in Exchange 2000 is the Win32 application programming interface (API) that allows drive letters to be mapped to the Web Storage System. HTTP access to the Web Storage System is better than HTTP access to data in earlier versions of Exchange. Database APIs such as OLE DB and ActiveX Data Objects (ADO) can also be used to access the Web Storage System.

Microsoft Office 2000 applications can browse Exchange public folder stores and mailbox stores, and open and save documents directly in Exchange.

Because of these new storage technologies and access methods, Exchange 2000 can host new types of corporate information. Exchange servers, with support for replication, online backup, and clustering, can serve as a central repository for both mail and Office documents. You can use Exchange 2000 support for multiple databases to set up additional public folder stores and experiment with these new capabilities without affecting your existing public folder stores.

Integrated Full–Text Search

The Web Storage System includes built–in indexing for high–speed, accurate, full–text searches; users can find content quickly and easily. Users see the same Outlook search interface that they are accustomed to, but the queries execute significantly faster than with earlier versions of Exchange. Instead of finding matches only in messages, users can find matches in documents that are attached to e–mail messages or saved directly in Exchange folders. All content in a mailbox or public folder store is indexed; this includes messages, documents, contacts, tasks, calendar items, and collaboration data. Index generation uses the same crawl technology as IIS and Microsoft SQL Server 7.0.

Scalability and Reliability

Exchange 2000 supports transaction logging for all databases. Support for multiple databases allows server storage capacity to grow while limiting the number of users who are affected if a single database must be restored.

New to Exchange 2000 is active–active clustering. Two–way clustering is supported with Windows 2000 Advanced Server. Four–way clustering may be available in a later update. This support enables storage groups to be allocated to cluster nodes and to fail over to other servers if necessary. Because all servers in the cluster are active and because no dedicated fail–over servers are needed, hardware is used more efficiently than in active–passive designs. Computer configurations do not need to be identical, which increases the number of options available to administrators.

The streaming capability of the Web Storage System and front–end and back–end architecture provide excellent performance if you use Internet protocols. The streaming capability means that MIME data received in SMTP and accessed using Internet protocols such as POP3 or IMAP4 does not undergo any intermediate transformation; this improves efficiency and eliminates potential conversion problems. The new front–end and back–end architecture allows administrators to dedicate one set of servers for client communication (front–end servers) and a separate set of servers for client data (back–end servers). The front-end servers can use Network Load Balancing and share a common name; the back-end servers can be in a server cluster. This allows administrators to scale particular parts of the system as needed. In addition, it ensures that SMTP failures and malicious attacks will not affect the messaging data or the directory.

Front–end and back–end architecture can provide firewall protection for Outlook Web Access. For organizations with two layers of firewall protection, front–end servers for Outlook Web Access can exist in the semi–trusted network; this provides a high level of control over information flow through the firewall.

Real–Time Collaboration

Although e–mail is accepted as a standard method of corporate communication, Exchange supports communication in real time that complements stored e–mail messages. Chat Service and Instant Messaging Service (available in Exchange 2000) allow people to exchange text messages in real–time. Exchange 2000 Conferencing Server (available as a separate product) provides the ability to schedule a data or videoconference from a messaging client.

Chat Service

With Chat Service, people can join live text discussion forums–-channels—thatare organized into virtual communities. Discussions on any number of topics can be ongoing or spontaneous, public or private. An Exchange server can host several chat communities, with each community serving a different interest.

Instant Messaging Service

Users of Instant Messaging Service can determine which subscribers to the service are online and available and then conduct private text conversations with them. Although widely used on the Internet, instant messaging is a newcomer to the corporate environment. It offers several key features:

- **Peer support** The ability to pose a quick question can greatly improve communication, particularly in widely dispersed teams. Not all questions are suitable for more formal e-mail communication.
- **Presence information** The Instant Messaging client monitors keyboard activity to determine if a user is at his or her computer. Presence information can be controlled manually; third-party developers can add support for additional devices that update presence information when the user is physically in front of the computer. This information can then be used, for example, by routing applications that determine which user to send work items to.

Exchange 2000 Conferencing Server

Exchange 2000 Conferencing Server, the new member of the Exchange product line, advances Microsoft's vision of "meetings without walls." You can purchase Exchange 2000 Conferencing Server as a product separate from Exchange 2000 Server. Exchange 2000 Conferencing Server allows knowledge workers to conduct data conferences, voice conferences, and video conferences across the intranet and Internet. It also gives administrators tools for administering back-end servers and provides support for bandwidth allocation, load balancing, and failover.

Conference Management Service

Exchange 2000 Conferencing Server includes a centralized reservation system that allows users to schedule and join meetings by using the Outlook calendar. Conference Management Service simplifies the management of conferencing services by providing a single interface for meetings of all types. In addition, Conference Management Service allows administrators to control the number of conferences that can be scheduled simultaneously.

Data Conferencing Provider

Data Conferencing Provider in Exchange 2000 Conferencing Server integrates with Conference Management Service; it supports the T.120 protocol and provides application sharing, whiteboard sharing, and chat capability. Data Conferencing Service broadcasts data to clients and manages conferences from beginning to end. Clients use Microsoft NetMeeting or other T.120 compatible programs to participate in data conferences.

Video Conferencing Provider

Video Conferencing Provider, which is available with Exchange 2000 Conferencing Server, uses Conference Management Service to provide multi-party, continuous presence video conferencing over an IP multicast. In addition, Video Conferencing Provider enables H.323 protocol clients that don't have multicast support to participate in video conferences.

Third-Party Support

Exchange 2000 Conferencing Server provides a standard mechanism for third-party developers to integrate their own conferencing technologies with Exchange 2000 scheduling, reservation, and access methods. Third parties can extend the capabilities already in the product; for example, they could provide an integrated H.323-protocol audio and video conference bridge for conferencing services or provide integration with telephony systems, such as voice conferencing servers.

Custom Solution Development

Exchange 2000 offers many new features for the developer, including XML support, enhanced Collaborative Data Objects (CDO), enhanced workflow, OLE DB and ADO support, server events, IIS and Active Server Pages (ASP) integration, and Web forms.

XML Support

Exchange 2000 provides complete, native support for XML, including all Exchange properties available in XML. It allows for the creation of unique data formats for specific applications and delivers structured data to the desktop for computation and presentation.

Enhanced CDO

Exchange 2000 includes significant enhancements to the Collaboration Data Objects (CDO) data model. CDO is a set of Component Object Model (COM) objects that are used for specifying business logic for workflow and other collaborative applications, developing Web-based applications, and accessing Active Directory. CDO can be used to create applications that take advantage of Exchange 2000 features such as messaging, scheduling, contact management, system management, Active Directory access, and workflow.

CDO 3.0 is built on OLE DB. In addition to providing access to Exchange data, it provides access to Internet services such as Lightweight Directory Access Protocol (LDAP) queries and MIME messages. By using CDO 3.0, administrators and developers can add capabilities to both the server and the Outlook client to suit their technology and business needs. For example, they can archive messages, enforce corporate policies, forward notifications to pagers, and manage distribution list traffic.

Other enhancements in CDO 3.0 include:

- The ability to create complex MIME multi–part messages.
- Support for S/MIME and X.509 v3 certificates.
- Expanded support for calendar and contact management using vCard, iCalendar, iTIP, iMIP, etc.
- Dual–interface and programmatic control using Microsoft Visual C++, Microsoft Visual Basic. Visual Basic Scripting Edition, Java Script, and Java.

Note that Exchange 5.5 custom applications run without modification on Exchange 2000 Server.

Server Events

Exchange 2000 provides a comprehensive server event model, which enables designers to create many types of applications that were previously impossible. Two kinds of events exist: transport events and Web Storage System events (both synchronous and asynchronous).

Transport events are used to track and customize the basic operations of Exchange 2000 while the system processes SMTP and NNTP requests. For example, an application could use transport events to attach a corporate disclaimer on all outgoing mail, rewrite the mail headers, or force all mail sent to a distribution list with more than 1,000 addresses to go through an approval process.

Web Storage System events, based on OLE DB events, occur when Exchange processes content in the Web Storage System; events are particularly useful for workflow and tracking applications. For example, an expense reporting application could initiate an accounts payable workflow process when a user saves an Excel spreadsheet into a particular Exchange folder.

Synchronous events are new to Exchange 2000. When a synchronous event is triggered, Exchange performs no further processing on an item until the event's associated business logic is executed. For example, a synchronous event could be associated with a folder that checks the validity of the digital signature of an expense report before the expense report is sent to management for approval.

Enhanced Workflow

Exchange 2000 includes CDO Workflow Objects—an enhanced library of workflow services based on simultaneous, synchronous events that provide a high–performance, reliable and secure engine for departmental and enterprise workflow and tracking applications. Exchange 2000 can drive a workflow tracking process such as document approval, expense reporting and payment, or purchase orders. CDO 3.0 and the Web Storage System event model provide a systematic workflow engine to control workflow state–transitions and to associate business logic actions with events. Business logic is programmed using Visual Basic; it can be created using standard tools.

OLE DB and ADO Support

The Web Storage System includes semi-structured databases of items that can be navigated by using familiar SQL syntax provided by OLE DB and ADO interfaces. OLE DB and ADO 2.5 are programming interfaces that provide a common means of accessing Microsoft BackOffice data, regardless of which server application the data is stored in. Exchange 2000 includes two OLE DB providers: a remote provider for server access from client applications such as Outlook 2000 and Office 2000; and a local provider implemented natively in Exchange for high-performance COM access from applications such as virus scanning programs and workflow engines.

OLE DB generalizes row-set access, query specification and execution, and data hierarchy navigation. In addition, application developers can use their experience with SQL tools such as queries. forms and reports, and they can combine SQL and Exchange data manipulation. Developers can also use ADO to navigate, query, filter, and sort Exchange Server data. This means that developers familiar with SQL applications can use the same expertise and tools to write applications that use data stored in the Web Storage System or that access both Exchange and SQL Server data.

IIS and ASP Integration

Exchange 2000 is a platform for high-performance Web-enabled applications, which provides unparalleled integration with Microsoft IIS and ASP technology. Powered by the new capabilities of the Web Storage System, Exchange 2000 can host Web sites and workflow applications that use either Outlook or a browser for data access and that can be created by familiar application design tools such as FrontPage 2000, Visual Basic, Visual Basic for Applications, and the Microsoft Visual Studio development system. Web projects can also take advantage of Web Storage System properties, events, business logic, and other features—all within a single repository.

Exchange 2000 significantly increases the functionality and performance of Web-enabled applications by rendering HTML and executing ASP scripts directly in the core process of Exchange 2000, through the use of XML and through other Internet standards. This enables developers to more closely integrate Web Storage System content, such as contacts, documents, meetings, tasks and workflow processing, into Web applications.

Web Forms

Web forms are browser–based forms that are stored in the Web Storage System and transmitted by Exchange in HTTP directly to the browser. For example, a folder could contain an expense report form or purchase order form that a user can access by typing a simple URL or selecting a bookmark in a Web browser. Web forms can be created with FrontPage 2000; they enable a user to create and modify items in a folder by using any browser that supports HTML 3.2. Internet Explorer 5 provides additional user interface features for forms, including drag–and–drop capability. This technology gives application developers a powerful, integrated tool for linking Web applications to the business logic, events, and metadata properties of the Microsoft Web Storage System.

You can use FrontPage 2000 to edit and manage Web applications hosted by Exchange 2000. For example, you can create Web–based, custom forms in FrontPage 2000 and have them hosted by Exchange. Using the dialog boxes in FrontPage, developers can open a Web page in an Exchange folder, edit it with FrontPage, and then run the application in a browser.

CHAPTER 2

Building the Project Plan

John Fisher, Senior Consultant, Microsoft

The project plan is essential to planning, designing, and deploying Microsoft Exchange 2000 Server. A thorough project plan provides logical phases that keep the team focused on the tasks involved. This chapter addresses items that you need to consider when creating your project plan.

Deploying Exchange 2000, whether as an upgrade from Microsoft Exchange version 5.5 or as a new installation, is a complex project that requires a plan to aid the people working on the project. To assist in the planning and deployment of a project, Microsoft Consulting Services and the Microsoft development groups have assembled a framework of best practices called the Microsoft Solutions Framework. Microsoft Solutions Framework is not a methodology, but a flexible series of concepts, models, and best practices that lay the foundation for planning, building, and deploying technology projects.

This chapter uses concepts from the Microsoft Solutions Framework Infrastructure Deployment Process Model, which is based on a team approach and milestone-driven schedules. The team that deploys Exchange can be an internal division of a company or an external consulting company that works with many companies. The scenario outlined in this chapter is based on an internal Information Technology (IT) team deploying Exchange 2000 for its own company. You can easily adapt the guidelines in this chapter for a team that consults for several companies and is not part of the company deploying Exchange 2000.

> **Note** Microsoft Solutions Framework is part of the Microsoft Enterprise Services Framework, a framework that targets different aspects of providing world-class information technology to companies. For more information about Enterprise Services Framework and its components, see thef Microsoft Web site at http://www.microsoft.com.

In This Chapter

Taking a Phased Approach

Phase I: Envisioning

Phase II: Planning

Phase III: Developing

Phase IV: Deploying

Taking a Phased Approach

When you are planning any large project, it is easy to get overwhelmed; dividing the project into phases helps keep the team on track and makes it easier to measure progress. Taking a phased approach to your project reduces risk and increases project success by ensuring that necessary tasks are completed in a logical order. Each phase of a project should have clear objectives and deliverables that, once completed and reviewed, allow the team to proceed to the next phase of the project.

This chapter discusses the four-phase model used with the Microsoft Solutions Framework Infrastructure Deployment Process. The phases are:

- **Phase 1: Envisioning** In this phase, you define the project vision and scope, you create the project team, you gather business requirements, and you identify project risks.
- **Phase 2: Planning** In this phase, you plan and design the solution. From a design perspective, you create detailed, low-level specifications for the project in the form of a functional specification. In addition, you establish detailed project plans in the form of the Master Project Plan and Master Project Schedule to help you complete the project.
- **Phase 3: Developing** In this phase, you develop the final system by using the design and planning information gathered in the planning phase. You finalize and tune systems in pilot projects, and create deployment plans.
- **Phase 4: Deploying** In this phase, you deploy the system. You take the project from the pilot stages in the last phase to full production status.

Assessing Project Risk

Assessing potential problems that might arise during your project—and devising a method to mitigate them—can reduce the negative impact they might have on your project. Risk management is a team effort, which continues throughout the life of the project; however, attempting to identify the risks early and mitigate them keeps the project on track and on schedule.

You can use the following five-step risk management process on your project:

1. **Identify the risk** This step allows the team to become aware of a risk and address it before it can cause harm to the project.
2. **Analyze the risk** This step involves analyzing the risk and converting it into information that the team can use to make decisions. The risk is assessed as to its probability of occurring, as well as its potential impact on the project. Using this information, the team can devise a plan for mitigation.
3. **Plan for risks** This step involves creating plans to turn the information that is collected into actions and decisions. The plan should involve developing actions to address the risks, prioritizing risk actions, and devising a metric or trigger point to alert the team that the mitigation plan must be put into action.

4. **Track the risks** Tracking the top ten risks is important; it allows the team to monitor them if mitigation steps were taken.

5. **Control risks** This step moves the risk management process into the daily activity of the project. If the team did a good job on the earlier steps, risk control should be easy by monitoring the metrics and triggers. Risk management remains a high-profile activity to aid in the success of your project. After a risk can no longer have an impact on the project, it should be retired.

Creating a table to assess risks (by estimating the probability of the problem arising and the impact it may have on the project) allows you to mitigate the risk. The mitigation may be a full plan or just an item or step to complete to eliminate or reduce the risk to the project. Although you create the table at the beginning of the project, it is constantly managed throughout the life of the project, and risks are added, removed, and retired as required.

Example

The following example can help clarify the risk assessment process.

A company is in the process of installing a new online training system for company-wide use, but there is a concern that it will not be in place in a timely fashion for the Exchange deployment. Based on the current schedule, the team has assigned this risk a high probability and a high impact on user satisfaction. As a result, the team tracks the risk as follows:

Risk Identifier: Online User Training.

Risk Statement: The training plan calls for all user training to be delivered by using the new online training system that is currently behind schedule for being installed. If this schedule continues to slip, the team will have no mechanism for delivering training material to users.

Risk Management Strategy: A member of the User Education team responsible for training monitors the installation status of the online training system on a weekly basis; if time allows, this person will offer assistance.

Risk Management Strategy Metrics: The Online Training System project schedule is no longer slipping or the team is actually close to getting back on schedule.

Action Items: Assign a team member to monitor and assist with the Online Training System and report back to the Exchange team weekly.

Due Date: Ongoing.

Personnel Assignment: Sandra from the User Education team.

Risk Contingency Strategy: Work with an outside vendor to obtain printed training materials rather than the online versions, and ensure that enough copies are available for the deployment.

Risk Contingency Strategy Metrics and Trigger Value: If the Online Training System slips more than an additional three weeks, or if 30 days prior to the first deployment, online training is not available, plan to order printed training material for user self study.

Identifying, analyzing, mitigating, and tracking the risks early in the project cycle allows the team to use the metrics and trigger point to easily identify when the risk needs to be addressed. Tracking the top 10 project risks in this manner eliminates the guesswork of risk management and aids in the success of your project.

Templates to Aid with Your Project

To assist in your program planning, several templates have been included on the *Microsoft Exchange* 2000 *Sewer Resource Kit* companion *CD.* You can use these templates to develop the various plans required for the project.

- Deployment Schedule Template
- Action Item Template
- Vision-Scope Template
- Risk Management Plan Template

Phase I: Envisioning

In the first phase of the project, you define the goals, limits, and structure of the project. The key deliverables in the envisioning phase are:

- Structure of the project team
- High-level project requirements
- Project vision, scope, and assumptions
- Conceptual high-level design

Note While in this phase of the project, it is especially important to identify and assess possible high-level project risks that you may encounter. For more information about risk assessment, see the section "Assessing Project Risk" earlier in this chapter.

Many project teams want to immediately begin the work of designing a new system; however, it is important to complete the envisioning phase first because this is when you assess customer requirements and create the final vision and scope of the solution. Without the final vision, a project can become unmanageable.

The envisioning phase concludes when your vision and scope document is complete and approved by the team and the project sponsor.

This section explains the major components of the envisioning phase.

Building the Project Team

The creation of a team allows you to distribute project responsibilities among team members. In the Microsoft Solutions Framework model, a team of peers plans and implements the project. The overall project responsibility lies with the entire team, but each member of the team is responsible for his or her appropriate functional area.

The different roles of an Exchange deployment team might include:

Table 2.1 Deployment team roles

Role	Responsibility
Product Manager	Works with the customer to gather business requirements, set objectives, and set the budget. This is not necessarily a technical role. The product manager is also responsible for project communications.
Program Manager	Assumes responsibility for the overall technical Exchange design and implementation.
Executive Sponsor	Provides support at a management level throughout the life of the project. As a high-level manager (usually a Director, Vice President, or above) this person is not necessarily a team member, but an external advisor to the team.
Project Sponsor	Reviews the progress of the project. This is an individual or small team of key stakeholders from the company. This person may be a director or manager to whom the technicians deliver the project report. The project sponsor is not necessarily a team member, but rather an external influence over the entire team.
Development/Engineering	Determines all technical configurations of the Exchange messaging system and interfaces, including the server, clients, and external connections. The development team consists of people functioning in different roles, such as: • Messaging Administrator • Network Administrator • Internet/Intranet Administrator • Desktop Administrator • Account Administrator • Security Administrator • Operations Engineer These professionals work together, on the central team or in sub-teams as required, to design the new system.

Table 2.1 Deployment team roles *(continued)*

Role	Responsibility
Test/Quality Assurance	Ensures that the designed systems comply with the functional specification and other corporate standards.
User Education	Ensures that the user education process and documents are completed, including all documentation and training for the project. Skill areas utilized on this team include: • End User Training • End User Technical Support • User Communications
Logistics Management	Determines the best way to deploy Exchange servers and transition systems to the operational groups.

In a larger project, you can assign roles to different sub-teams that will focus on the functional area. In a smaller project, you can assign a single role to an individual team member or you can combine roles, so that one member handles multiple roles.

Often an IT group that undertakes a project on its own may choose to contract an external technical consultant to advise the project team. This resource can provide occasional consulting services to the team across all teams; Microsoft Consulting Services as well as Microsoft Certified Solutions Providers can provide assistance with this role.

All team members should be proficient with Exchange 2000 and Microsoft Windows 2000 so that the team can share a base level of understanding. This training is important so that the team understands the effort required to add functionality to the systems. For additional information about training classes and authorized training centers, see the Microsoft Web site at http://www.microsoft.com.

Building the Project Structure

You must set up and communicate the operating structure during the envisioning phase of the project. In the next phase, you will add details to the structure. Depending on the culture at your company, different degrees of formality may be acceptable, but at a minimum, you should hold regular scheduled project meetings with an agenda. Consider the use of a project Web site to store all pertinent project documents, schedules, and status reports, so that all team members, as well as the customer. can access the latest information.

Gathering High-Level Requirements

Gathering the requirements of your new Exchange environment is important because these requirements make up the infrastructure or framework of the system design; they can be categorized as follows:

- **Business requirements** The requirements, such as service level objectives, that the business demands of the new Exchange messaging cystems. For example, you must decide if Exchange will be the building block for future business applications.
- **User requirements** Specific features that the users have requested.
- **IT requirements** Requirements from IT management as well as systems administrators, which outline the structure and operation of the Exchange environment.

When you collect requirements from the different groups, you might want to prioritize requirements or create a list of requested features. By having the customers prioritize their requests, you enable your team to phase in features if the project timeline is aggressive.

Defining the Project Vision

A high-level project vision or mission statement helps the team focus on its tasks when designing and deploying the new Exchange system. The vision statement describes what the team envisions as the end result or outcome of the project.

For example, a fictional company called Northwind Traders has 10,000 employees in fifteen locations worldwide. The company has five different business units, each maintaining its own messaging system, ranging from Microsoft Exchange Server 5.5 to an old mainframe text-based mail system. To reduce IT costs and improve communications, Northwind Traders' IT group has decided to replace the old messaging systems and install a new company-wide standard that uses Exchange 2000. The project team has created the following vision statement for its project:

> Northwind Traders will replace its five existing divisional-based messaging systems with an enterprise-wide messaging system that uses Microsoft Exchange 2000 Server. This will enable the five business units to interoperate seamlessly while adding user functionality and reducing administrative and operational costs. The new system will be deployed to each of the company's locations and all employees worldwide. Deployment will commence within three months and be complete within 18 months.

Defining the Project Scope

Whereas the vision statement defines a final vision of the project, the project scope defines the functionality that you can reasonably implement within the time allowed, given the project variables. Project variables include:

- **Resources** People and money.
- **Schedule** The time allowed for the project.
- **Features** The specific features that are included in the final solution.

  For example, the new messaging system at Northwind Traders will:

  - Be designed for 10,000 users and accommodate Northwind Traders' 7 percent expected growth per year, for the next seven years.
  - Replace the five previous messaging systems: Microsoft Exchange, Lotus Notes, Lotus cc:Mail, UNIX Post Office Protocol (POP) servers, and SNA Distribution System (SNADS).
  - Coexist with the five legacy messaging systems until migration is complete.
  - Be a highly available system that fulfills the requirements specified by the company's business units' Service Level Agreement.
  - Provide the capability for real-time employee communications to facilitate idea sharing across the company.
  - Leverage the newly installed Windows 2000 infrastructure.
  - Provide an internally hosted data conferencing service to eliminate the need to contract this service from an outside provider.

Defining the Project Assumptions

Assumptions should be identified as early as possible in the process and listed in the project's vision and scope document. The addition of items outside the project's scope should also be listed in the assumptions section to clearly identify functionality that will not be delivered as part of the project.

For example, the following are assumptions for an Exchange 2000 deployment:

- The Windows 2000 infrastructure project must be complete enough to use Active Directory directory service for this project.
- The resource gap will be filled with contract labor to meet the aggressive project schedule.
- The migration team is not responsible for mail older than 30 days.

Creating a Conceptual Design

Often when planning a project, you can develop a conceptual design based on your vision statement and the requirements that you have gathered. The conceptual design would be a high-level drawing of your optimal solution and would be included in your vision and scope document. It is important to not go into too much detail at this phase because the detailed design occurs in the planning phase when you create the functional specification.

Phase II: Planning

The second phase of a project is the planning phase. This phase plays an important role in the smooth deployment of your Exchange environment. In the planning phase, the team completes high-level concepts in the envisioning phase and begins the detailed planning and engineering tasks of the system. Planning, also called the engineering phase of a project, addresses three main points through the following deliverables:

- Functional Specification: Details the final solution that the team delivers.
- Master Project Plan: Details how the system is designed, tested, and deployed.
- Project Schedule: Details when the remainder of the project begins and when it will be completed.

To cover the What, How, and the When of the project, the team must complete several major tasks during the planning phase. General planning tasks help the team understand the current environment and plan for the new project's environment. The following are general planning tasks:

- Gather information
- Design the functional specifications
- Perform proof-of-concept testing
- Identify resources
- Create the project plan
- Build the project schedule

Specific planning tasks must include addressing the following issues:

- It is critical to understand the existing Windows infrastructure, whether it is Windows NT version 4.0 or Windows 2000. If Windows NT 4.0 is currently being used, then the team must understand how and when Windows 2000 will be deployed in the company.
- If Exchange 5.5 is currently deployed, the team must plan an approach for interoperability, integration, and user migration to Exchange.
- A clearly defined owner of Active Directory, both for planning and operational purposes, must be identified. If you do not plan for this issue, you might later face political and operational barriers. If Exchange is already deployed, Exchange administrators may be accustomed to owning the directory. When Active Directory is introduced into the environment, Windows NT administrators and not Exchange administrators might own Active Directory.

The planning phase is complete when the team agrees to the functional specifications and the program plan.

Gathering Information

When designing the new messaging system, it is critical that the team has an accurate picture of the existing environment. The following table illustrates the important information that you must obtain.

Table 2.2 Information about the existing environment

Information to Gather	Description
Organizational structure and location data	How the business is organized and operates (for example, divisions vs. regions). This often dictates message flow. The number of office locations and users at each location is also important.
Network infrastructure	The company network infrastructure (including LAN and WAN topologies) and utilization. It is important to determine how the company handles remote access.
Messaging and directory structure	What the current messaging topology looks like and how it operates. The existence of a Windows 2000 Active Directory and whether it needs to interact with other directories within the company is also important.
Desktop and server environments	The types of systems that users are currently using. It is important to determine if existing messaging hardware can be re-deployed for use with Exchange 2000.
Standards	Whether there are existing corporate standards, such as naming and operational procedures, that impact the design of the system.
Functional requirements	These requirements expand upon the high-level user requirements. It is a good idea to conduct user and IT surveys to determine user and IT requirements.

The Functional Specification

After the teams have gathered the appropriate information, you use this information along with the requirements, vision and scope documents, and high-level design to start creating a functional specification. The functional specification must provide enough detail to allow the various team members to begin work in their areas, such as engineering, testing, training, and deployment.

As the functional specification evolves, it may become clear that there is not enough time to complete the system with all of the desired functionality, particularly the desirable but non-essential items. To deliver a solution that meets the project schedule, consider limiting the functionality included in the release to the functionality that is required. Using the concept of versions, you can add functionality to the system in future projects. When deciding on which features to move to a next version, always refer to the vision and scope, as well as user requirements. These items, along with estimated costs and scheduling information, aid in the decision.

The functional specification can be considered a contract between the company and the project team on what will be delivered for a final solution. At the end of the planning phase, the team and the customer approve a draft version of the functional specification. This provides a mechanism for setting expectations for everyone included in the project. In the next phase, after the design is tested, the functional specification is frozen. To ensure timely delivery, changes to the design are no longer acceptable.

Proof-of-Concept Testing

During the design and engineering process, you might need to validate some engineering designs by performing a prototype or a proof-of-concept test. These tests are not intended to test the entire solution, but to aid in the engineering process and verify that a design is actually achievable. Often teams choose to prototype a new technology, such as conferencing services or connectors to external systems. This helps the team better understand product features and aids in the design of the system.

For example, given that Northwind Traders' IT group already has experience with Exchange 5.5, they will prototype some of the new or updated features and connectors that are required for the project. These features include:

- Exchange Instant Messaging
- Exchange Conferencing Server
- Third party e-mail connectors, such as Lotus Notes and Lotus cc:Mail.
- Multiple storage groups

Identifying Resources

In the next step of the planning phase, you identify resources such as staff, hardware, software, and tools required for the project. These resources help in the creation of the project plan and schedule.

When identifying staffing resources, consider the following questions:

- Do you need to fill the team members' regular positions for the duration of the project?
- Are the team members properly trained for designing and deploying Exchange 2000 and Windows 2000?
- Do you need to hire consultants to augment your team?

When identifying hardware resources, consider the following questions:

- Do you have hardware available for a development and lab environment?
- Can you use existing production messaging servers to support Exchange 2000?

When identifying software resources, consider the following questions:

- Can you upgrade your existing software licenses to Exchange 2000?
- Which version of Exchange 2000 do you need to purchase?
- Do you need to obtain Windows 2000 licenses for Exchange users or have the licenses already been purchased?

When identifying tools resources, consider the following questions:

- Do you need to obtain a tool to aid in the migration of old mail?
- Do you require additional monitoring tools for your environment?

Creating Required Project Plans

A project has multiple project plans, each listing the tasks that an individual team member performs. All of these plans form a Master Project Plan. You can consider the Master Project Plan the blueprint of your project, which details the planning of your project and helps ensure a smooth deployment.

The major sections of a Master Project Plan might include:

- Product and Program Management
- Architecture/Engineering Document
- Test Plan
- User Pilot Plan
- Security Plan
- Budget
- Training Plan
- Communications Plan
- Deployment/Migration Plan
- Logistics/Operations Plan
- Risk Management Plan

Building the Project Schedule

After completing a draft of the functional specification and master project plan, the team builds the schedule, based on the work outlined in the specification and master project plan. The team also decides on a release date for the final system. The project schedule includes the details of the four phases and the work required from the different team members. The team may need to reassess the schedule and modify the functional specification to achieve an acceptable release date for all involved.

It is important for the team to develop the master schedule and agree to it, because the team members will be held accountable for the schedule and will be measuring their progress against it. Creating detailed schedules with various milestones allows team members to measure their progress and quickly identify any issues that may impact the project schedule.

A sample Exchange 2000 schedule template created in Microsoft Project 2000 is included on the ***Microsoji Exchange*** 2000 ***Sewer Resource Kit*** companion CD. This template is organized around the four Microsoft Solutions Framework phases and includes the various tasks for each phase.

Phase III: Developing

The third phase of the project involves testing and verifying the designs and plans completed in the planning phase. The teams set up the equipment to be used in production, and follow the plans and designs to create the final systems or a portion of the final environment. A series of verification and pre-pilot tests occur on the system and the phase is complete when the user pilot testing has been performed and the systems are ready for the first production use.

The major steps in the developing phase are:

- Validate the design and project plan
- Build the system
- Complete a pre-pilot test
- Complete a user pilot test

Validating the Designs

You must validate the designs in a lab environment that contains hardware identical to the hardware that will be used in the production environment. You should verify the design aspects of the system according to the functional specifications and then test and verify installation procedures. It is important to approach the testing in a systematic way by using the test plan. For information about how and what to test, see "Setting Up a Test Environment" in this book.

Testing first begins with the individual components or servers to verify that the components will perform as configured according to the specifications. The tests identify possible problems that may arise in production because of a setting or configuration. After the individual components and servers perform as expected, you should test them at the systems level. In systems level testing, you test multiple servers or components to ensure proper operation.

> **Note** Remember that these tests are still in a stable environment and final testing will be complete with the pre-pilot and user pilot testing. The verification tests merely verify the design and ensure that it is ready for the building of production systems.

Building the System

After you have successfully validated your designs in a lab environment, you can begin building the system with production hardware on the production network. At this point, the systems are almost ready for users. After the system is in place, you can begin the final testing phase.

Pre-Pilot Tests

The production hardware and systems are now in place and ready for a user pilot test. However, it is a good idea to perform a pre-pilot test of your system with a small group of technical users or project team members to verify the system operations. Having a group of technical users test the system before you make it available to mainstream users allows you to catch any bugs that slipped through the validation and testing. The chapter "Piloting Exchange 2000," will help you determine which features should be tested and how to test them. It is important to remember to follow the same procedures in the pre-pilot test as the procedures that will be used during the actual user migration, because the pre-pilot should test the procedures as much as it does the systems.

After receiving feedback from the pre-pilot participants, you can tune systems if required; then the system is ready for the final step: the user pilot.

User Pilot Test

After completing the pre-pilot test and reconfiguring the system, you should perform a user pilot test, which is the final step prior to actual deployment. The user pilot test should be performed with a group of users that are representative of all users. The size of the pilot group varies depending on the size of your deployment. Pilot users are usually volunteers because they must accept less reliability from a system not yet ready to be used by the whole company. Support, operations, and training should be in place as if it were a production deployment.

You should include the following items in the user pilot test:

- Proper communications of the migration or rollout
- User training
- User documentation
- User install process
- User support
- Public folder, e-mail alias, or Web site setup for pilot feedback
- Pilot evaluation

After the user pilot test is complete, you must perform an evaluation of the user pilot test. You must make a determination of whether the systems and procedures are ready to be deployed into production. It may be a good idea to partner with the pilot users to determine an acceptable criteria level for the new system. After approval by the project team and pilot users, the system is ready to be deployed.

It may be necessary, after receiving user pilot test feedback, to make modifications in migration procedures, user training, or the system itself. Remember to plan for this prior to deployment.

> **Caution** Scope creep occurs when the customer of the new system attempts to add functionality to the system after the requirements are frozen. For example, often when the system undergoes pilot testing, users request features that may not have been included in the original requirements. There is a tendency to attempt to add this requested functionality at the last minute. However, it is best to deal with these requests in later versions of the system so that they do not impact the original scope of the project.

Phase IV: Deploying

The fourth and final phase of the project is deployment. In this phase, all users are migrated to the new messaging system according to the deployment or migration plan. Training is completed for users, and the operations team monitors the new system. After this phase is completed, the nortnal production and operations groups take control of the systems.

During this phase of the project, the deployment teams take control and the other teams are used to support the deployment as needed. Engineering and testing may be required for problem resolution, but the deployment teams are responsible for the majority of the work in this phase.

Migrating Users

Company-wide technologies such as messaging backbones, gateways to other systems, and Internet connections, should be deployed first. Server clusters in large data centers may also be considered core technology. After the core technologies are installed, you can identify servers to install in smaller locations.

Migrate users to the new system according to the migration plan that was finalized in the developing phase. Train users on the new system and train support staff to support the new system and users.

For example, Northwind Traders will deploy its main data center location first because it hosts the connectors to the Lotus Notes, Lotus cc:mail, and mainframe systems as well as the connection to the Internet for Simple Mail Transfer Protocol (SMTP). After the main data center is complete and stabilized, the deployment team deploys the local mailbox servers in the smaller regional offices.

Handing Off the Project to Production

When all of the user mailboxes have been migrated or deployed on the new system, hand the system off to the operations team. The systems are now considered production systems. The message operations team, or a similar group, monitors the systems and ensures proper operation.

Completing the Project

After the new system is in place and all of the deliverables are complete, the project team formally closes the project. The team establishes agreement that the project has been completed and deployed as specified in the project plan. The project team then conducts a project evaluation, in which it reviews areas that could be improved in future projects, and documents these findings for future project teams.

Proper project planning is a critical component to the overall success of your project and should be an integral part of your deployment of Exchange 2000. It is important to use a model that allows your team to be flexible but maintain a systematic approach that keeps the team on schedule while accomplishing the project goals. Many projects fail or run into difficulties partway into the deployment because the proper planning was not completed. By using these planning steps or a similar systematic approach to program planning, you will be on your way to a successful rollout of Exchange 2000.

PART 2

Planning for Exchange 2000 and Active Directory

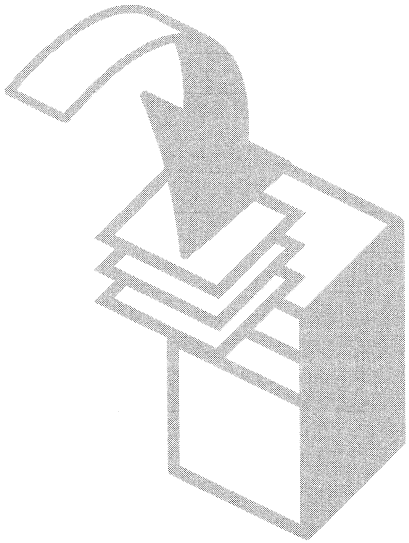
In This Part

The Exchange 2000 Environment

Active Directory Design

Active Directory Integration and Replication

Deployment Strategies



CHAPTER 3

The Exchange 2000 Environment

Paul Sebben, Managing Consultant, E-Sync Networks, Inc.

Microsoft Exchange 2000 Server provides e-mail and scheduling functionality, online forms, Web access to mailboxes and folders, and the tools for custom collaboration and messaging-service applications. To enhance the reliability, scalability, and performance of the overall system, Exchange relies on software, hardware, and networking components such as Microsoft Windows 2000, the Active Directory directory service, and the Domain Name System (DNS). Exchange 2000 also uses a single Web Storage System that can handle messages, discussion threads, Microsoft Office documents, Web documents, and voice-mail messages, among others. All of these pieces work together to create the Exchange 2000 environment.

Successfully deploying Exchange is like solving a puzzle. You must consider the size and shape of the puzzle pieces before you can assemble them effectively. If a piece of the puzzle is missing, you cannot complete the puzzle. The purpose of this chapter is to identify the key pieces of the puzzle that will assist you in successfully implementing Exchange 2000.

This chapter describes how to identify the critical factors required for a successful deployment of Exchange by reviewing company organizational structure, existing messaging environment, network infrastructure, and business requirements. This chapter describes each key topic area required for the deployment, along with its impact on Exchange 2000. At the end of each section is a series of questions that will assist you in the architecture phase of Exchange 2000. As you answer the questions, you should identify whether your existing infrastructure will accommodate Exchange 2000, whether you need to upgrade critical areas, or whether you must build from scratch those critical areas that do not already exist.

In This Chapter

Company Structure

Business Requirements

Current Network Infrastructure

Windows NT 4.0 Domain Structure

Active Directory

Exchange Architectural Review

Server Configuration

Public Folders

Organizational Forms

Conclusion

Company Structure

Windows 2000 and Exchange 2000 can have a tremendous influence on a company's structure. In many companies, the DNS, network, Windows, directory service, and Exchange teams typically work independently of one another. To successfully deploy Windows 2000 and Exchange 2000, you must create a team that can work as a cohesive unit to implement this technology. Executive management support is imperative in deploying this technology because the teams involved often work for different individuals within an information technology (IT) department. You must identify ownership of each technology area, including the following areas:

- Name resolution services, such as WINS and DNS
- Microsoft Windows NT and Windows 2000 Server
- Active Directory implementation
- The public key infrastructure
- Network infrastructure

After you determine what the existing teams are and how they work together, you should create a team or an environment in which all the teams work closely together to deploy Exchange 2000.

You must also examine the IT management structure. In a centralized environment, the IT department determines the standards that are deployed within the corporation. The decision-making process for deploying new technology is typically easier because fewer individuals are involved. In a completely decentralized environment, standards between divisions often do not exist. Therefore, more individuals must be involved in the decision-making, project-planning, and deployment processes. Questions to consider include:

- Is this environment centralized, decentralized, or both?
- In a combination centralized and decentralized environment, is a centralized IT office responsible for cross-divisional directions and standards?
- Is the company global?
- Does everyone speak the same language, or is a translator required to assist during meetings and with user communication?

Answering these questions helps you identify how messaging and collaboration standards are developed. In addition to examining the IT structure, you should examine the routine operational tasks on the servers. These include:

- User and mailbox creation
- Service pack upgrades
- Message flow
- Backup and restore

Examining these daily operations helps you determine the areas that will be most affected when you deploy Exchange 2000. You should be aware of the level of service these servers provide to the company so that you can continue to meet service expectations.

Business Requirements

Assessing the company's business needs is a key strategy for gaining approval to deploy Exchange 2000. Companies do not perform upgrades just so that they can use the latest version of a product. A business justification for investing in the new technology is required. This justification can be divided into the following areas:

- Organizational requirements
- User requirements
- Administrative requirements

Use these guidelines to research ways in which Exchange 2000 can help you achieve company goals for better collaboration, provide users with better access to resources, and offer easier administration of the messaging and collaboration environment.

Organizational Requirements

E-mail has revolutionized the way organizations communicate, and is now considered a critical application by most companies. A company's organizational requirements define their corporate standards for e-mail use. These often include policies dealing with legal implications, such as e-mail retention, as well as standards describing the proper use of e-mail. They also define user service levels, and other initiatives, including ways to reduce the total cost of ownership.

Your first step is to determine your company's organizational requirements. After you are aware of the existing requirements, you can use Exchange 2000 to find new ways to meet and possibly even improve organizational requirements.

User Requirements

User requirements define the features needed to allow users to perform job responsibilities in a more efficient and productive manner. To gather these requirements, the project team must communicate with users. Because user requirements are not typically described in technical terms, the project team's responsibility is to map the user requirements to technical requirements. The product will then be evaluated to make sure that the user needs are being met. The following table gives some examples of user requirements and their technical requirement counterparts:

Table 3.1 User and technical requirements

User Requirement	Technical Requirement
To collaborate with other users interactively.	Establish Instant Messenger services.
To easily publish and modify HTML documents that are stored in public folders by using Microsoft FrontPage or Microsoft Word.	Implement Exchange 2000, the Web Storage System, and the Installable File System.
To quickly locate messages and documents stored in public folders.	Enable full-text indexing in mailbox stores and public folder stores.

To gather these requirements, you should determine your company's user requirements. After you are aware of existing requirements, you can use Exchange 2000 to find new ways to maintain or possibly even improve user productivity.

Administrative Requirements

Administrative requirements define the configuration and management structure for Exchange 2000. This includes:

- The use of policies for managing mailboxes.
- The ability to separate user accounts into smaller databases for faster backup and restore operations.
- The ability to define a group of users who manage only the routing topology.

After you determine your company's administrative requirements, you can use Exchange 2000 to find new ways to meet and possibly even improve the administration of Exchange 2000 and Exchange user accounts.

Future Planning

It is important to forecast the future business needs of your user environment. This includes planning for new technologies. Microsoft Exchange 2000 enhances the user environment by integrating several new technologies, including:

- Instant Messaging
- Online conferencing
- Chat services

Some of the important planning information that you should gather includes:

- Is the company currently involved in acquisitions and divestitures?
- Will the current user base remain the same, or will it increase or decrease? Will growth be at a single location or spread out across the organization?
- Will the company add any new locations within the next year?
- Is the company planning to deploy instant messaging, online conferencing, chat services, or unified messaging?

The answers to these questions will be useful when planning for server hardware, clustering, administration, and routing needs of the new Exchange 2000 environment.

Current Network Infrastructure

To effectively plan the Exchange 2000 installation, you must examine the network infrastructure. Understanding this infrastructure assists you in determining how routing groups should be configured; it also allows you to build redundancy into your Exchange organization. In addition, it allows you to troubleshoot problems with mail flow and helps you determine if there are sufficient global catalog servers near your Outlook users.

Depending upon the size of the company, gathering this information may require contacting multiple individuals from various geographical areas. For a small company, usually there is only one contact person. For medium to large multinational companies, you may need to contact several teams to obtain the necessary information.

The Windows 2000 deployment team may have already gathered some or all of the necessary information. Be sure to check with the Windows 2000 team to see what information they have on hand.

WAN Environment

To provide a high-level view of the network, you should create a network topology diagram. Divide the diagram by geographical area, and include connections to all sites regardless of the connection type. Include in the diagram the speed, circuit type, and Committed Information Rate (CIR) for each WAN link. Figure 3.1 illustrates a sample network diagram.

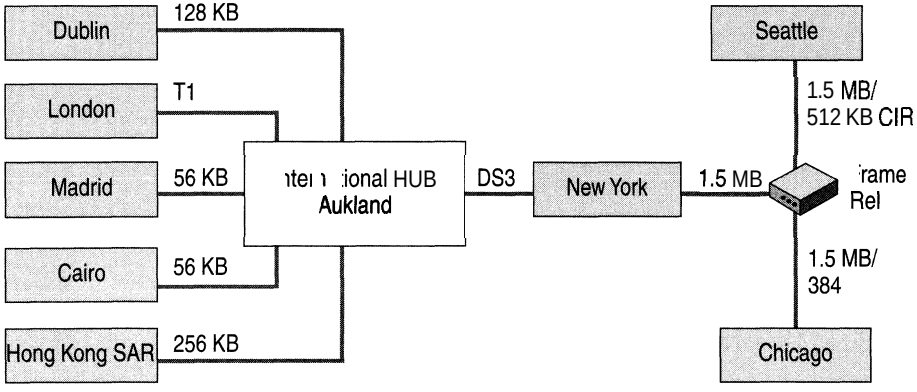


Figure 3.1 Sample wide area network diagram

You should also create a table for each WAN link. Constructing this table helps you determine how the WAN is currently functioning and how routing groups should be configured. This table should include the following information:

- The originating location for the circuit
- The destination location for the circuit
- The speed of the circuit
- The circuit type, for example, T1, *56* Kbps, *256* Kbps, frame relay, virtual private network (VPN), and so on
- The Committed Information Rate (CIR)
- Whether there is an alternate path to this LAN
- The current available incoming bandwidth
- The current available outgoing bandwidth
- Whether this is an Internet VPN connection
- The protocols being used on this link
- The Internet Protocol (IP) information regarding this circuit

LAN Environment

You must also examine and document the LAN connections within each geographical location. If a graphical representation of the LAN does not already exist, you should create one. Include in this diagram information detailing how the subnets within the facility are connected. The diagram's complexity will vary depending upon the size of the location.

The LAN diagram you create should contain the following information:

- The type of network being used at this location, for example: Ethernet or token ring
- Whether there are any established firewalls, and if so, how they are configured
- The speed of the LAN, for example: asynchronous transfer mode (ATM), 10 MB, 16 MB, 100 MB, and so on
- The protocols that are being used, for example: IP, Internetwork Packet Exchange/Sequence Packet Exchange (IPX/SPX), and so on
- The subnets that are being used within this location
- The subnets that are being defined but not used
- How the subnet is connected, for example: switches, routers, and so on
- Whether local WINS servers are defined for this location, and if so, what the IP addresses are for these servers
- Whether there are local DNS servers defined for this location, and if so, what the IP addresses are for these servers
- The protocols used by clients to access the Exchange servers
- Whether there is a static IP addressing scheme used for this location
- The remote access strategy for this location
- Whether users connect to this location or another location for access, if remote access is required
- Whether Outlook Web Access is deployed at this location

Constructing the LAN diagram helps you identify how each area is currently connected, if new routing groups should be created, or if DNS servers exist.

Domain Name System

Previous versions of Exchange used WINS as the preferred naming service; Domain Name System (DNS) was required only for Internet, Simple Mail Transfer Protocol (SMTP), and Outlook Web Access. It is recommend that Exchange 2000 servers be placed in a DNS zone that is on a Windows 2000 DNS server. Clients query DNS to locate the nearest global catalog server, and Exchange servers use DNS to locate other Exchange servers. Therefore, a solid DNS implementation is required for the successful deployment of Exchange 2000, and an examination of the corporate DNS structure is required. Typical questions to consider include:

- Is DNS currently deployed internally? If so, under which operating system is DNS running?
- Which version of DNS is running?
- Which are the primary and secondary DNS servers?
- Has DNS been integrated with Active Directory?
- What are the DNS namespaces that are currently deployed within the organization?
- Are the clients configured to use DNS?

Answering these questions helps you identify how DNS is currently used. Because DNS is recommended for deploying Exchange 2000, if DNS is not being used, you should plan to move from WINS to DNS.

If you choose to deploy Windows 2000 and Exchange 2000, you should be running DNS bind version 8.1.2 or later because it supports SRV (service) records (RFC 2052) and dynamic updates (RFC 2136).

Windows NT 4.0 Domain Structure

Exchange 2000 is tightly integrated with the Windows 2000 operating system. Before upgrading to Exchange 2000, you must add your Exchange servers to a Windows 2000 domain and to an Active Directory forest. Exchange users must also be authenticated in the Windows 2000 Active Directory forest before connecting to Exchange 2000.

You must understand the existing domain structure before migrating the server to Windows 2000. Because the Windows 2000 domain model is tightly integrated with Exchange 2000 security, the Windows and Exchange teams should collaborate to minimize the impact on users when the directory migrates from Windows NT version 4.0 domains to Active Directory. You should create a global diagram and a domain diagram to illustrate the current environment. You will use these diagrams in planning the migration from Windows NT 4.0 to Windows 2000.

The global diagram should include:

- All domains, including trust relationships
- Whether trusts are one-way or two-way
- When and how each domain is scheduled for migration to Active Directory
- A description of plans for domain consolidation, if any

The domain diagram should show:

- Where the primary domain controllers are located for each domain
- Where the backup domain controllers are located for each domain

After you examine the domain structures, you can define the Active Directory migration path.

Active Directory

In previous versions of Exchange, the directory service was installed on each Exchange server in an organization. The directory service was responsible for storing the Exchange organization's configuration information. The most substantial change in Exchange 2000 is the removal of the directory service from Exchange. Active Directory is now a core component of the Windows 2000 operating system; all Exchange objects are stored in Active Directory.

Many of the technologies in Microsoft Exchange version 5.5 have been incorporated in Active Directory in Windows 2000. Active Directory provides the structure and functionality to manage user accounts, desktops, the network, and applications from a single location.

This section focuses on the Active Directory components that are required to effectively design an Exchange 2000 organization. These components include forests, sites, schemas, global catalogs, and global catalog servers. For a detailed explanation of Active Directory, see the *Microsoft Windows* 2000 *Sewer Resource* **Kit** *Deployment Guide.*

Forests

The number of Active Directory forests in an IT infrastructure directly impacts Exchange 2000 organizations. There is a one-to-one relationship between the number of Active Directory forests and Exchange 2000 organizations. The Exchange global address list can display users from only a single forest.

Sites

Windows 2000 sites are different from the sites in previous versions of Exchange. In Exchange 2000, routing groups and administration groups have replaced sites.

A Windows 2000 site is a group of servers that communicates over a permanent high-bandwidth connection and can be subdivided into IP subnets or IP ranges. Windows 2000 sites are important to Exchange 2000 because you use Windows 2000 sites to determine the location of the nearest domain controller that can validate logon credentials. In addition, Exchange 2000 stores its configuration information in Active Directory and relies on information in Active Directory to generate the global address list.

Schema

The Active Directory schema is an extensible, programmable catalog of all objects, attributes, and classes that exist in Active Directory. The configuration partition of Active Directory must contain the Exchange 2000 configuration data. This requires that Exchange extend the schema to include Exchange 2000 specific objects, classes, and attributes. Extending the schema is a forest-wide action that replicates new objects to every domain controller and global catalog in the forest.

Extending the schema requires planning to ensure that the process does not over-burden the network and affect user productivity. Exchange extends the schema the first time that you install Exchange 2000 Server. In some cases, it may be desirable to extend the Active Directory schema prior to installing or upgrading to Exchange 2000. You can accomplish this by running the Exchange ForestPrep utility.

You need to determine if the Active Directory schema has already been extended for Exchange 2000.

Global Catalog and Global Catalog Servers

The global catalog contains a complete replica of all objects in the domain and a partial replica of all objects in the forest. The global catalog is the component of Active Directory that Exchange 2000 clients query to obtain the global address list. The proximity of the global catalog server to the client directly affects the performance of the client. At least one global catalog server should reside in each Windows 2000 site.

By default, the global catalog does not replicate all attributes that were included in previous versions of Exchange, although additional attributes can be configured to replicate and display in the global catalog. For example, the attributes for first name, last name, and SMTP address are tagged for global catalog replication by default. However, the attributes for mailing address fax number, home telephone number, and custom attributes one through ten are not.

If you need to display a non-default attribute in the global address list, you must modify it for replication to the global catalog. During messaging evaluation, you should examine all attributes to determine if modifications are required. Modifying the replicated attributes triggers replication to the global catalog. It is important to plan and ensure that business processes are not affected when this happens.

Questions to consider include:

- Has the Active Directory schema been updated for the Exchange 2000 installation?
- What attributes need to be replicated to the global catalog?
- Is there at least one global catalog server located at each Windows 2000 site?

Answering these questions helps you determine the scope of the necessary preparation work for the schema, the number of attributes that you must replicate, and the availability of the global catalog. You should plan to have a global catalog server in as many areas as you need to maintain service expectations.

Active Directory Connector

The Active Directory Connector (ADC) allows for replication between Active Directory and an Exchange 5.5 directory. With ADC, you can control the direction in which the replication information flows between a master server and other servers in the organization or site. Active Directory uses a connection agreement to control replication between an Exchange 5.5 recipient container and a Windows 2000 domain controller. Questions to consider include:

- Has ADC been deployed?
- If ADC has been deployed, what connection agreements have been defined?

If ADC has not been deployed, you must plan time for this before deploying Exchange 2000. If ADC has already been deployed, review any existing connection agreements to determine if you must create new ones.

Exchange Architectural Review

To maximize benefits and reduce risk during the deployment of Exchange 2000, a thorough understanding of the existing messaging environment is very important.

You should examine the existing organization and site design, and pay special attention to:

- Mail routing
- Server configuration
- Client configuration
- Business needs

These are the critical areas for any messaging and collaboration system. If you focus on these areas and determine your existing design, you can provide basic messaging functionality before adding the more advanced features available in Exchange 2000.

Overall Exchange Design

The concept of organizations and sites has changed for Exchange 2000. The organization no longer uses an independent directory that resides on every Exchange server. Rather, it has been integrated with Active Directory and is bounded by the Windows 2000 forest in which it resides. Administrative groups and routing groups have now replaced Exchange 5.5 sites. Administrative groups are collections of Exchange 2000 servers that share a common administrative model. For example, consider an organization that has 100 Exchange servers. If the same group of administrators manages ten of those servers, you can place those ten servers in an administrative group. A routing group is a group of Exchange 2000 servers that are connected by a high-speed link. Servers within a routing group communicate directly with each other by using SMTP.

With Exchange 5.5, administrators commonly set up a different SMTP address for each site. The default in Exchange 2000 is one recipient policy with one SMTP address that covers all users in the entire organization. If needed, you can have multiple recipient policies so that different SMTP addresses apply to different people. Exchange 2000 enables you to have one SMTP domain for all users.

You should divide the existing Exchange organization into components and examine it with the new administrative and routing group model in mind. Exchange 2000 Setup prompts you to specify the administrative group to which the server will belong. After a server is installed in an administrative group, it cannot be moved. It is essential that you establish an administrative group before you install Exchange 2000.

The existing organization can be divided into the following components:

- Exchange 5.5 site design
- Backbone
- Internet connectivity

If you do not already have a diagram that reflects the current structure for each of these components, you should create one. This both aids in the design of the new Exchange 2000 organization and provides troubleshooting assistance in the existing organization. To assist in the creation of these diagrams, Microsoft has developed Microsoft Exchange Server Topology Diagramming Tool (EMap.exe). You can download this tool from the Exchange Web site at http:Nwww.microsoft.com/exchange.

You should also consider whether multiple Exchange organizations exist within the company. If multiple organizations do exist, you need to determine if it would be beneficial to combine them into a single organization or Active Directory forest. When determining whether to separate or combine multiple organizations, you should consider whether the organizations were created because of mergers or acquisitions, and how the directories between the ditferent organizations are connected.

Address Book Views

Address book views provide an alternate view of the global address list. They are defined at an organizational level and automatically replicated to all servers. In Exchange 2000, the address book views are called address lists and are built using Lightweight Directory Access Protocol (LDAP) queries. Consider the following questions:

- Are there any address book views defined?
- Has each address book view's configuration been documented?

Answering these questions helps you deploy Exchange 2000 without changing how address books appear to users. You may also find new ways to make address books more useful for users.

Exchange Site Design

Exchange 5.5 sites were created based on administrative and network connectivity needs. All servers within a site communicate with each other directly. and sites can span multiple geographic locations, provided sufficient bandwidth is available. In addition, user administration is easier within a site. Within a site, servers are usually categorized by their function. These categories include mailbox servers, connector servers, and Outlook Web Access servers. In small organizations, connectors and mailboxes may be kept on the same server.

Questions to consider include:

- What is the site addressing for each Exchange 5.5 site?
- What is the default SMTP format used for addressing user mailboxes?
- What are the defaults for mailbox display names and aliases?
- What recipient containers are used to generate the offline address book?

Reviewing this information allows you to upgrade Exchange 5.5 sites to Exchange 2000 routing groups, and administrative groups. It can also help you plan how Exchange sites will work with Exchange 2000 organizations.

Mailbox and Public Folder Servers

Mailbox and public folder servers contain only mailboxes and public folders. In theory, there is no limit on the number of mailboxes that can reside on an Exchange 5.5 server. In reality, the number of users is limited by several factors, including the amount of RAM, the processor speed, and the recovery time required by a critical failure. Exchange 2000 allows multiple mailbox stores to coexist on a single server. The ability to start and stop each mailbox store independently allows a single mailbox store to be restored without affecting users on a different mailbox store.

When you analyze the existing server configuration, you should examine the sizes of the mailbox and public folder stores. Consider the following questions:

- What is the default storage limit on mailboxes?
- What is the retention policy for deleted items?
- What is the distribution list strategy?
- Who is responsible for creating distribution lists?
- Are there automated tools used to create distribution lists?
- Which protocols are enabled for mail?
- How many mailboxes on a server are related to service level agreements with customers?
- Is there a need to provide VIP Mailboxes?
- Are there multiple companies within your Exchange organization? If so, is there a need to recover the companies separately?

If large mailbox stores exist, splitting them into several smaller ones should be considered. In the past, service level agreements have often required the deployment of multiple servers to ensure recovery within the service level agreement time constraint. Multiple storage groups and databases allow you to consolidate these servers into a larger server with multiple mailbox stores and databases.

Bridgehead Servers

In earlier versions of Exchange, servers that have one or more connectors connecting two Exchange sites, or connecting Exchange to the Internet are called bridgehead servers. Bridgehead servers are responsible for sending mail and public folder information between sites, and for replicating the Exchange directory between sites. In Exchange 2000, bridgehead servers still exist, but directory information is now replicated by using the global catalog. You should create a table with information about each bridgehead server connector. Questions to consider include:

- Are TCP or TP4 transport stacks being used?
- What type of connector is installed (for example, SMTP, X.400, Site Connector, and so on)? Document the configuration of the connector.
- What is the site or server destination for this connector?
- What is the cost for this connector?
- Are there any downstream sites below this connector?
- How often does this connector send mail?

- Is there an alternate route defined for this server?
- Is Exchange directory replication configured?
- What is the Exchange directory replication schedule for each site?
- Are there message limits on the connector? If so, document the message restrictions

If Exchange uses TP4, you must plan to convert to other transport stacks. Answering these questions helps you determine how bridgehead servers are currently used and how they are affected when Exchange 2000 is deployed.

Connector Servers

Connector servers interact with other servers. Some connectors that were supported in Exchange *5.5* are not supported in Exchange 2000. However, Exchange *5.5* sites that include servers with connectors provided with Exchange *5.5* or earlier can be used to transfer messages to Exchange 2000. Exchange 2000 supports the Lotus cc:mail, Lotus Notes, Microsoft Mail, and Novell GroupWise connectors. You should consider the following questions:

- Are there any Exchange *5.5* connectors installed, and if so, which ones?
- Is there documentation for each type of connector? For example: MS Mail, cc:mail, Lotus Notes, Professional Office System (PROFS), Systems Network Architecture Distribution System (SNADS), and so on.
- Is there documentation about the configuration of each connector? If not, document the connector's configuration.

Your need for Exchange *5.5* connectors influences whether you upgrade all servers to Exchange 2000 or leave some Exchange *5.5* servers in place to manage those connectors.

Backbone

A backbone diagram illustrates how the entire Exchange environment is connected. It includes all sites in the Exchange organization and the connectors that link them. It also includes connection costs and network connection bandwidth between sites.

Internet Mail

Internet mail makes up a significant percentage of message flow within an organization. It is essential that you document Internet message traffic before your upgrade to Exchange 2000. In earlier versions of Exchange, configuration information for outbound message content was stored on each Internet Mail Service. In Exchange 2000, content configuration is now managed for each organization.

Create a diagram that shows inbound and outbound message traffic to and from the Internet. Include information about firewalls and virus scanning. In addition, answer the following questions:

- Is there a firewall between the Internet Mail Service server and the Internet?
- If so, does the server queue messages or pass them through the firewall?
- Are there any size restrictions for messages at the firewall?
- Who maintains the external mail exchanger (MX) records for the company?
- What MX records are registered?
- Are the MX records configured in a load-balancing configuration?
- Does the Internet Mail Service send SMTP mail to any internal systems? If so, document the name of the system; the configuration information, such as IP address, DNS name, and so on; the business purpose; and the contact information.
- Is the Internet Mail Service configured for re-routing? If so, document the configuration.
- Are inbound and outbound messages scanned for viruses? Who is the anti-virus software vendor and how is the software configured? How often are virus patterns updated?

Answering these questions helps you plan for upgrading to Exchange 2000 and securing mail that is exchanged over the Internet. When you upgrade to Exchange 2000, configuration information that existed for each Internet Mail Service is not upgraded and now applies to an entire organization. You must migrate all of your Internet Mail Service configuration information to the Exchange System Manager before you add any SMTP connectors from Exchange 2000 servers to the Internet.

Disaster Recove Planning

When an Exchange server experiences a critical failure, a company's disaster recovery plan is tested. The ability to execute this plan quickly and effectively is crucial. Disaster recovery plans vary widely between companies. Some companies have no plan, others an untested plan, and some a fully tested plan. A fully tested plan that includes backup validation helps guarantee success and minimize downtime in the event of an emergency. Questions about disaster recovery should include:

- Is there a disaster recovery plan, and where is it located?
- Who is on the disaster recovery team?
- Has the plan been tested?
- Is there spare hardware used for this purpose, and if so, where is it located?
- What is the backup schedule?
- Are the backups working? Have they been validated?
- What is your Product Support Services account number and telephone number?

Having a fully tested plan in place before you upgrade to Windows 2000 or Exchange 2000 is important. By answering these questions, you can determine bow long it currently takes to recover from a disaster, and fix any problems before introducing more advanced software into the recovery plan.

Exchange Monitoring and System Performance

Before you upgrade to Exchange 2000, you must have a thorough understanding of the health of your existing environment. You should examine event logs and collect system performance data.

System performance data gathered before the Windows 2000 or Exchange 2000 upgrades helps determine whether to upgrade or replace system hardware. After you install Exchange, you can gather new data to determine if the system is performing as desired and to confirm that decisions involving hardware were correct. Consider the following questions:

- What is the daily and weekly message volume for each server?
- How many mailboxes are installed?
- What is the level of processor activity?
- What is the disk throughput (read/write activity)?
- How much memory is being used?

Event logs should be scanned for errors regarding Active Directory, Microsoft Web Storage System, or message transfer agents. Certain errors, such as a -1018 error, could prevent a successful upgrade to Exchange 2000.

Server Configuration

Exchange 2000 contains several new features that affect server configuration. You must document the existing configuration before you design the Exchange 2000 organization. Examine every server in the Exchange organization.

Hardware

You must examine each server to determine if it meets the system requirements for Windows 2000 and Exchange 2000. Hardware requirements include processing speed, RAM, and hard disk size. You should also examine page files. Typical configuration questions include:

- How many CPUs are installed?
- What is the speed of each CPU?
- How much RAM is installed?
- Where are the page files located and how large are they?
- Are the transaction logs on a separate volume from the databases?

You should also consider the configuration of disk subsystems by asking yourself the following questions:

- What type of disk controller is used?
- How much memory does the controller have?
- What type of redundant array of independent disks (RAID) array is deployed?
- How many spindles are present?
- In which format is each disk partition?

If your current hardware does not meet the system requirements for Windows 2000 and Exchange 2000, you need to upgrade your hardware before deploying Exchange 2000.

Clustering

E-mail is a critical activity for most companies, and users expect the messaging system to always be available. Unfortunately, hardware failures occur. and operating systems periodically need upgrades and service packs. This results in scheduled and unscheduled downtime. In certain environments, downtime is unacceptable. Clustering can provide users with a highly available system by reducing or eliminating single points of failure. **A** cluster is a group of independent servers that share a common disk subsystem, act as a single computer, and have a common network name and IP address. The server's hardware configuration is identical for each server in the cluster.

In a cluster, if one server is designated as the active node, the active node has exclusive control of the shared disk and all cluster resources. The other node in the cluster is in a passive (standby) state waiting for the active node to fail. When failure occurs, the passive node takes control and becomes the active node. This type of clustering enables an administrator to apply upgrades during the day without affecting users. The administrator can apply the upgrades on the passive node, and then change this node to the active node. This process is transparent to users and has no impact on them. Exchange mailbox servers are typically used in a cluster.

Consider the following questions:

- Is clustering installed in your existing Exchange organization?
- Is there a plan in place for implementing clustering?

If you are not using clustering in your current installation, consider the benefits of doing so. Windows 2000 Advanced Server supports two-node active-active clustering. In active-active clustering, each node shares a portion of the processing load. If one node fails, the other node takes over the user load. This action remains transparent to users and can reduce the cost of implementing a cluster because servers are no longer idle until a disaster occurs.

Windows 2000

Several key Exchange 5.5 components have been integrated into the Windows 2000 operating system. These components include the directory and transports such as HTTP, Network News Transport Protocol (NNTP), and SMTP. Because of this integration, you must install Windows 2000 before you install Exchange 2000.

To minimize the impact of an upgrade, you can upgrade the operating system of Exchange 5.5 member servers in a Windows NT 4.0 domain to Windows 2000; these servers can still be members of the Windows NT 4.0 domain. You should consider the following questions:

- What version of Windows NT is currently installed?
- What service packs have been installed?
- Where are the page files located and what size are they?

By identifying the operating system configuration, you can determine any necessary upgrades that must be performed before Exchange 2000 is deployed.

Exchange Software Version

Before you upgrade to Windows 2000, you must install Exchange 5.5 Service Pack 3 (SP 3). The Active Directory Connector requires Exchange 5.5 (SP 1) or later when operating in a Windows NT 4.0 domain.

Third-Party Software

Increasingly, Exchange servers are running software in addition to Exchange. Third-party software, including backup software, monitoring software, virus checkers, and paging and fax software are often installed on the same server as Exchange. This presents challenges for Exchange administrators. All third-party software on the server must be checked to determine if it is Exchange 2000–compliant. Questions to ask are:

- What third-party software is installed?
- Does the software support Windows 2000 and Exchange 2000?
- Can the software be upgraded to Exchange 2000 or does it require removal and reinstallation after Active Directory is installed?

You may need to remove and re-install any software that interacts with the directory, such as paging and fax software, after you install Active Directory. New versions of the software may also be required.

Server Names

Server naming conventions usually describe the location and purpose of Exchange servers. By looking at the server names, you can easily identify the location and function of a server. Naming conventions are especially important in a Windows 2000 environment. Windows 2000 is tightly integrated with DNS, and the underscore character is no longer supported in server names. If this character is contained in any existing server names, you must change these names prior to migrating to Windows 2000. Migration will fail if the server name is not changed.

Client Access

Many companies are unable to deploy new client software simultaneously with the deployment of Exchange 2000. Backwards compatibility with previous versions of Exchange is essential. All clients you use with earlier versions of Exchange, including Microsoft Outlook 97 and Microsoft Outlook 98, are compatible with Exchange 2000. In Exchange 2000, the global catalog provides the global address list to users. Microsoft Outlook 2000 is currently the only client that can access this directly. Exchange 2000 relays address list information to allow previous versions of Exchange access to the global address list.

Outlook 2000 is the premier messaging client for Exchange, and it is tightly integrated with Active Directory and Microsoft Office applications. It provides collaboration application support, filtered offline synchronization, and a local copy of the calendar.

In addition, Outlook Web Access has been improved significantly for Exchange 2000. Outlook Web Access now looks and feels like a native desktop application when accessed with Microsoft Internet Explorer 5. It features drag-and-drop functionality, popup menus, toolbars, and rich-text HTML text editing. Existing Outlook Web Access environments should be examined to see how they were used previously. Questions include:

- How many user logons occur per day?
- How many messages are being read and sent per day?
- What is the average session time?

You must also gather information regarding the client configuration, including:

- What messaging client software has been deployed? For example, Exchange, Outlook 97, Outlook 98, Outlook 2000, Outlook Express, Outlook Web Access, and so on.
- Which version of each client is deployed?
- Is there a plan for upgrading to Outlook 2000, and if so, when is this scheduled to occur?
- Which Web browser is installed? Which version?

Answering these questions helps you determine the client load on Exchange and the extent to which Outlook Web Access will be used. You may also want to plan to upgrade any Exchange clients or Web browsers that are heavily used.

Public Folders

A public folder stores messages or information that users in your organization can share. Public folders can contain different types of information, ranging from custom forms to Internet content stored in its native format. You can create as many top-level public folder hierarchies as you need, but new public folder hierarchies that are not created in the default hierarchy are not accessible by MAPI clients such as Outlook. Each public folder hierarchy is represented by its own public folder store and can be replicated to other servers in the organization. After being replicated, if one server becomes unavailable, a client can use an alternate server to access public folder content. In these situations, Exchange affinities use a routing group that is configured to refer the client to another server. This referral allows the client to access content when a specific server is not available or known. Public folder content can be indexed and is accessible to Exchange clients, Web browsers, and custom applications.

Public Folder Hierarchy

The public folder hierarchy in an Exchange organization helps organize public folders into collections of information that are easy to browse. Typically, the public folder hierarchy is organized to reflect a company's internal organizational structure. A well-defined public folder structure is essential for any Exchange server implementation, allowing administrative tasks to be delegated. Tasks such as adding permissions, adding folders, and removing folders can be performed at the user level or delegated to a group administrator. You should examine your existing public folder hierarchy before you deploy Exchange 2000. Gather the following information about the existing public folder hierarchy:

- Are public folders a highly used service?
- Are the top-level folders logically organized in an understandable way?
- Who can create top-level folders?
- What permission levels do the users of the folders possess?
- Do administrators spend too much time administering public folders?

Answering these questions helps you determine how public folders are used within your company and the resources that will be necessary to continue to provide or to improve the use of public folders in Exchange 2000.

Public Folder Affinity

Public folder affinity allows users from one Exchange site to access public folders on another Exchange site. For example, most of the users might reside on one large central site. This site might store all of the public folders for the company, perhaps on a dedicated public folder server. If remote Exchange sites need limited access to the public folders, you can use affinity to allow access to the folders at the central site.

In Exchange 2000, when a client uses an alternate server to access public folder content, Exchange uses routing groups to calculate the closest available server. The closest available server, when no replica is available in the same routing group, is determined by a cost property set for the routing group connector. The cost for each routing group connector is stored in a single cost database that is shared with e-mail routing calculations. The redundant cost tables maintained in earlier versions of Exchange are eliminated in Exchange 2000. Routing group connectors also determine whether you can use another routing group's replicas. When you upgrade to Exchange 2000, public folder affinities are not upgraded.

Public folder affinity applies to all connectors that connect routing groups. These connectors include:

- **Routing Group connector** This connector always connects two routing groups.
- **SMTP connector** This connector can connect two organizations or connect an organization to a third-party system.
- **X.400 connector** This connector can connect two organizations or connect an organization to a third-party system.

You should understand certain aspects of public folder affinity before you begin an Exchange 2000 upgrade. Consider the following issues:

- Is public folder affinity in use?
- Which sites participate in affinity?
- Is affinity necessary, or would replicas be a better solution?
- What are the costs associated between affinities?
- Create a diagram showing all affinities within the organization, and determine where affinity has been established between sites.

Understanding how affinities are currently used is critical. In Exchange 5.5, site affinities are not transitive. For example, if you set up an affinity between site 1 and site 2, and between site 2 and site **3,** you do not automatically get affinity between site 1 and site 3. In Exchange 2000, affinities are transitive; in the above situation, affinity exists between site 1 and site 3. If one routing group is connected to another, all servers receive public folder referrals. However, you can configure specific routing group connectors to deny public folder referrals. You may want to consider implementing new public folder affinities after Exchange 2000 is installed.

Replication

Public folder replication is similar to public folder affinity, except that the sites that participate in replication actually house the folders, rather than connecting to a remote site. Public folder replication is determined based on two factors: user utilization and available bandwidth. If user access to public folders is not a priority, you can conserve bandwidth by opting not to replicate public folders to remote servers. In such cases, you can use an affinity instead. Sites where public folders are highly used should contain a replica of all the folders.

Until Exchange 5.5, administrators in one site could replicate folders from other sites without the permission of remote administrators. Exchange 5.5 introduced the "Limit Administrative Access to Home Site" option that prevented this from occurring. Examine the following aspects of public folder replication before deploying Exchange 2000:

- Whether public folder replication is being used
- When replication occurs
- Which folders are being replicated to which sites
- Whether affinity is sufficient for public folder access
- Whether your Exchange site contains unwanted folders from remote sites

You should carefully examine each site's public folder replicas, and limit the amount of replication to a minimum.

Full Content Indexing

In Exchange 5.5, you can use Microsoft Site Server to create an index of public folder content that users can search. Exchange 2000 Server creates its own index of public folder and mailbox content that users can search. You can enable indexing for any database so that content can be selectively indexed. Before planning your indexing policy, you should determine if you are already using Site Server to index your public folders. If you are using Site Server, you must set indexing schedules and properties in the databases in Exchange 2000 so that indexes continue to be created.

Organizational Forms

In Exchange 2000, each organization has a forms library to store forms commonly accessed by all users in a company. Forms enable users to enter and view information. For example, a standard supply request form can be stored in a form library. You can create forms from Exchange System Manager in Microsoft Management Console (MMC) or from an Exchange client. Consider the following questions:

- What custom forms are currently being used? For which platforms, such as Outlook and Outlook Web Access, were they developed?
- What workflow applications are currently being used?

If forms are currently being used in your company, you should ensure that they are available after Exchange 2000 is deployed.

Conclusion

This chapter identified the critical areas that must exist for a successful deployment of Exchange. Areas such as the organizational structure, existing messaging environment, network infrastructure, and business requirements were discussed in terms of key deployment issues and their impact on your messaging and collaboration environment. If you answered the questions at the end of each section, you will have the information you need to complete the architecture phase of Exchange 2000. If you found that your existing infrastructure does not accommodate Exchange 2000, you should have the necessary information to upgrade critical areas, or build the critical areas that do not yet exist.

CHAPTER 4

Active Directory Design

Sybil Wood, Technical Writer, Volt Technical Services
Markus Vilcinskas, Program Manager, Microsoft

The Microsoft Windows 2000 Active Directory directory service extends the features of Windows-based directory services and adds entirely new features. This chapter describes the Active Directory namespace, including the forest and tree domain structure, organizational units, global catalog, trust relationships, and replication. This chapter also discusses how Microsoft Exchange 2000 Server uses the Active Directory namespace components, and how Domain Name System (DNS) influences the namespace.

Because the design of Active Directory has a tremendous effect on Exchange 2000 design and performance, information from the *Microsoft Windows* 2000 *Sewer Resource Kit* is included in this chapter along with information about Exchange-specific functionality and requirements.

In This Chapter

Active Directory Overview

Active Directory Logical Structure

Active Directory Logical Components

Active Directory Overview

Active Directory is secure, distributed, partitioned, and replicated. It is designed to work well in any size installation, from a single server with a few hundred objects to thousands of servers with millions of objects. Active Directory adds many new features that make it easy to navigate and manage large amounts of information, saving time for both administrators and users. These features include a new domain model and hierarchical namespace.

The advantages of using one directory rather than several directories include centralized management, centralized administration, and a single replication mechanism. Most corporations that use multiple directories attempt to unify their directories either through complex replication algorithms or through the use of unified directories such as Active Directory. Exchange 2000 removes the complexity of creating a unified directory because it uses Active Directory as its sole directory for services and data storage.

When you install Windows 2000 Server, Internet protocol stacks such as Simple Mail Transfer Protocol (SMTP) and Network News Transfer Protocol (NNTP) are configured as part of the operating system. The operating system and other components use these stacks; for example, it is possible to replicate Active Directory information using SMTP rather than remote procedure call (RPC) communications. The Exchange 2000 installation extends these stacks with additional command verbs and advanced routing components to provide all of the functionality required for an enterprise-class messaging and collaboration system.

Namespaces in Your Company

A namespace is a set of unique names for resources or items used in a shared computing environment. Namespaces can be found in almost all systems (network operating systems, software tools, directories, DNS, messaging systems, database systems, and so on). The names in a namespace can be resolved to the objects they represent. For DNS, the namespace is the vertical or hierarchical structure of the domain name tree. For example, each domain label, such as *host1* or example, used in a fully qualified domain name (FQDN), such as host1.example.microsoft.tld, indicates a branch in the domain namespace tree.

Active Directory is primarily a namespace, as is any directory service. A namespace is any bounded area in which a given name can be resolved. Name resolution is the process of translating a name into some object or information that the name represents. Active Directory forms a namespace in which the name of an object in the directory can be resolved to the object itself.

For Active Directory, namespace corresponds to the DNS namespace structure and it resolves Active Directory object names. The domain is a collection of resources on the network that includes one or more domain controllers. The domain namespace consists of domains, trees, and forests, and defines the boundary of the domain. All Active Directory domains are identified by a DNS-style naming convention and a name that is compatible with network basic input/output system (NetBIOS). For example:

DNS-style domain name: seattle.microsoft.tld

NetBIOS name: SEATTLE

In general, the NetBIOS name will be similar to the first naming component of the DNS-style name; however, NetBIOS restricts the type and number of characters in the name and subsequently the two names may look different from one another. Upon creation of the domain, the Active Directory administrator can configure both the DNS-style and NetBIOS domain names.

Note In the initial release of Windows 2000 Server, Active Directory domains cannot be renamed.

A forest is a collection of one or more Windows 2000 domains that share a common schema, configuration, and global catalog and are linked with two-way transitive trusts.

Planning the DNS Namespace

The resolution of names through DNS is central to Windows 2000 Server operation. Without proper name resolution, users cannot locate resources on the network. It is critical that the design of the DNS namespace be created with Active Directory in mind and that the larger namespace that exists on the Internet not conflict with a company's internal namespace.

The recommended approach to DNS design is to design the Active Directory environment first and then support that design with the DNS structure. However, in some cases, the DNS namespace might already be in place. In such a configuration, the Active Directory environment should be designed independently and then implemented either as a totally separate namespace or as a subdomain of the existing namespace. If the namespace you choose already exists on the Internet, it might cause name resolution problems for internal clients. Consider the following when you plan the DNS namespace:

- Identify the DNS namespace that you will be using for your domain. Identify the name that your company has registered for use on the Internet (for example, *company*.tld). If your company does not have a registered name, but you will be connected to the Internet, you may want to register a name on the Internet. If you choose not to register a name, make sure that you choose a name that is unique on the Internet.

  Note Throughout this chapter, the root domain name *.tld* is used to distinguish fictitious or sample domain names from actual domain names. Typically, the root domain name is .com, .edu, .org, and so on, rather than .tld.

- Use different internal and external namespaces. Internally, you could use *company*.tld or a subdomain of the external name such as *corp.company*.tld. The subdomain structure could be useful if you already have an existing DNS namespace. Different locations or departments can be named with different subdomains such as *nameone.corp.company*.tld or *nametwo.corp.company*.tld to ease administration.

- Make Active Directory child domains immediately subordinate to their parent domains in the DNS namespace. You can choose to create subdomains for departments or locations within your company. For example, *leveltwo.levelone.corp.company*.tld

- Separate internal and external names on separate servers. External servers should include only those names that you want to be visible to the Internet. Internal servers should contain names that are for internal use. You can set your internal DNS servers to forward requests that they cannot resolve to external servers for resolution. Different types of clients require different kinds of name resolution. Web proxy clients, for example, do not require external name resolution because the proxy server does this on their behalf.

  Overlapping internal and external namespaces are not recommended. In most cases, the end result of this configuration is that computers cannot locate needed resources because they receive incorrect Internet Protocol (IP) addresses from DNS. This is particularly a concern when network address translation is involved and the external IP address is in an unreachable range for internal clients.

- When you run the Active Directory Installation Wizard (Dcpromo.exe), you can configure a DNS server on the local computer and configure the forward lookup zones. Active Directory Installation Wizard examines the TCP/IP configuration on the computer and determines whether the computer is configured to use any DNS servers. If so, Active Directory Installation Wizard queries the root servers. If the computer is not configured to use any DNS servers, Active Directory Installation Wizard queries the Internet root servers. If it cannot contact any root servers, it configures the local computer as a root server and creates the "." (root domain) zone.

  Make sure that root servers are not created unintentionally. Active Directory Installation Wizard can create root servers, resulting in internal clients being able to reach external clients or parent domains. If the "." zone exists, a root server has been created. It may be necessary to remove this for proper name resolution to work.

Defining the Namespace Architecture

When designing an Active Directory namespace architecture, you should consider the following design criteria:

- Replication traffic requirements
- The ability to accommodate company restructuring without expensive domain changes
- The ability to evolve the design as company needs change

Domain Structure

The basic unit in Active Directory is the domain. An Active Directory domain allows you to group and administer a collection of resources on the network. The domain boundary defines the namespace and includes one or more domain controllers.

A single domain can span multiple physical locations or Windows 2000 Server sites. Unlike Windows NT 3.*x* and Windows NT 4.0, which used primary domain controllers (also known as PDCs) and backup domain controllers (also known as BDCs), Active Directory uses a multimaster peer controller model called the domain controller. All domain controllers that are authoritative for a given domain can receive changes directly and propagate those changes. This allows replication to occur between sites within a domain, even if a domain controller is down.

Domain Controller

A server that can authenticate users for a domain is called a domain controller. There must be at least one domain controller in each domain. Each domain controller holds a complete replica of the domain naming context for the domain to which it belongs, and a complete replica of the configuration and schema naming contexts for the forest. You can promote a member server to a domain controller using Active Directory Installation Wizard (Dcpromo.exe). In addition, you can demote a domain controller back to a member server by using Active Directory Installation Wizard.

The First Domain

The first domain within an Active Directory forest plays an extremely important role. In fact, the name of the entire forest is based on the DNS name of the first domain that is installed in the forest. There are two important factors to take into consideration when defining the first domain:

1. The first domain can never be removed from the forest.
2. Other domains cannot be represented above the first domain in the domain tree. If the first domain is named cat.microsoft.tld, it is not possible to create microsoft.tld afterwards; however, other domain trees such as msnbc.tld can be installed in the forest afterwards. The key to remembering this is that you cannot create domains with names that are part of the first domain, but you can create domains that form other trees.

The design strategy for most large companies is to create the first domain as a placeholder for the rest of the forest. The first domain contains only domain controller computer accounts. If you plan to create this domain and store it in a separate tree from the rest of the company, you should still have the domain name registered on the Internet, and it should not have an arbitrary name, such as firstdomain.tld.

Building this placeholder domain ensures that companies that have decentralized information technology departments can bring their domains online when they are ready without having to wait for other departments to catch up. This prevents departments from creating their own forests.

Additional Domain Controllers

The number of domain controllers you will create for a given domain is driven by two factors: fault tolerance requirements and load distribution requirements.

For each domain, use the following guidelines to determine if more domain controllers are necessary:

- Always create at least two domain controllers

  Even for small domains with small user populations, create at least two domain controllers so that there is no single point of failure for the domain.

- For each Windows 2000 site that contains a single domain controller, decide if you trust the WAN for fail-over.

  Should the single domain controller fail, clients in the Windows 2000 site can be serviced by domain controllers for that domain that are located in other sites. If network connectivity is unreliable or intermittently available, you might not want to trust the network to handle fail-over. In that case, place a second domain controller for that domain into the Windows 2000 site.

- Place additional domain controllers for a domain into a Windows 2000 site to handle the client workload.

The number of clients that a particular server can handle depends on the workload characteristics and the hardware configuration of the server. Client computers randomly select from the available domain controllers in a Windows 2000 site to distribute client load evenly.

Domain Mode

An Active Directory domain can be in either mixed mode or native mode. In mixed mode, the domain is restricted to limitations (such as 40,000 objects) imposed by the Windows NT 4.0 domain model. However, it is possible to place a Windows 2000 domain controller in a Windows NT 4.0 domain. For this to happen, the first Windows 2000 domain controller takes on the role of the Windows NT 4.0 primary domain controller. While a domain is in mixed mode, the Windows 2000 domain controller functions as a Windows NT 4.0 primary domain controller with the scaling constraints of a Windows NT 4.0 primary domain controller. Windows NT 4.0 backup domain controllers see the Windows 2000 domain controller as the primary domain controller. When multiple Windows 2000 domain controllers are present in a mixed-mode domain, you can dictate, through the administration interface, which Windows 2000 domain controller functions as the primary domain controller.

> **Note** Microsoft Exchange 2000 Server also runs in either mixed mode or native mode. Although the concept of mode is similar in both Windows 2000 Server and Exchange 2000, the mode in Windows 2000 Server is different than the mode in Exchange 2000.

To gain the scalability benefits and functionality of an Active Directory domain, you must switch the Windows 2000 domain to native mode. Switching to native mode is irreversible and requires that all domain controllers be upgraded to Windows 2000. Native mode operation allows Active Directory to scale up to millions of objects and overcomes the constraints of the previous Security Accounts Manager (SAM). A domain in native mode allows for rich group creation and nesting, which is advantageous for Exchange 2000. A native-mode domain cannot include Windows NT 4.0 domain controllers, but can include Windows NT 4.0 member servers. Thus, Windows NT 4.0 clients do not need to be upgraded to belong to a native-mode domain.

Having your domains in native mode not only provides the operating system with additional scalability, but also simplifies the Exchange 2000 installation process, because it is possible to create and use universal security groups for public folder access.

Domain Tree

Domains are represented in a parent and child hierarchy known as a domain tree. A domain tree is a set of one or more Windows 2000 domains with contiguous names. Standard DNS domain names are used to represent the tree structure (for example, exchange.microsoft.com). Because microsoft.com does not have a parent domain, it is considered the tree root domain. The child domains of microsoft.com are exchange.microsoft.com and nwtraders.rnicrosoft.com. A grandchild domain of microsoft.com is jp.nwtraders.microsoft.com. These domain names are contiguous because each name has only one label different from the name of the domain above it in the domain hierarchy. Figure 4.1 presents a single tree with a contiguous namespace.

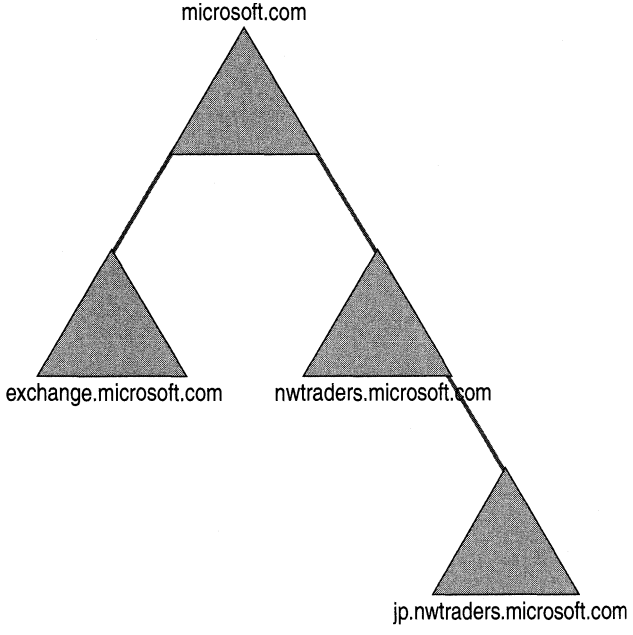


Figure 4.1 Single tree with four domains

Domains within the forest that do not have the same hierarchical domain name are in a different domain tree. When different domain trees are in a forest, the tree root domains are not contiguous. *Disjoint namespace* is the phrase used to describe the relationship between different domain trees within the forest. A multiple-tree forest is shown in Figure 4.2.

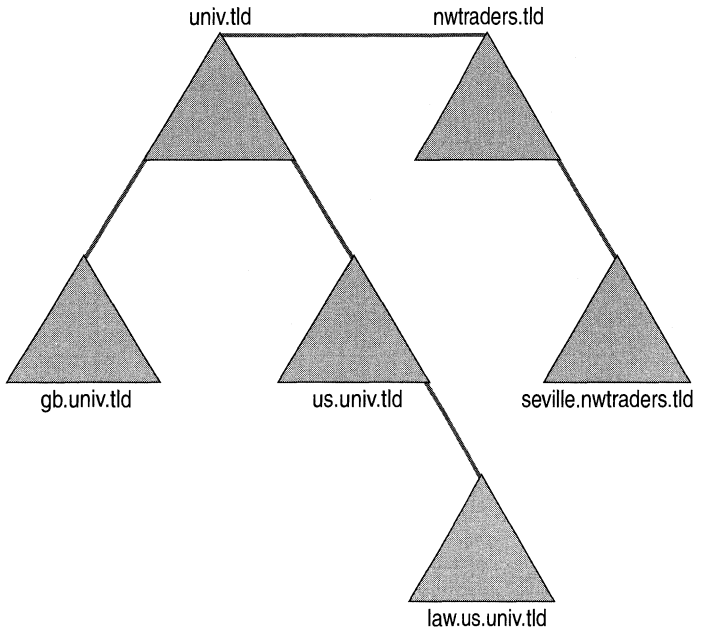


Figure 4.2 Forest with multiple trees

Domain Name System Overview

The DNS namespace is a distributed database organized as a hierarchical tree. Each node of this tree is called a domain and has a name. The first node in this tree is the root node and has an empty string as a name. Each node can have at least one child, which is also known as a subdomain.

Resource Records

Each DNS database consists of a set of resource records, which have one of the following formats:

```
[TTL] [class] type RDATA
[class] [TTL] type RDATA
```

The basic goal of a resource record is to associate a value with an entry. In addition, each resource record has a type identifier. For example, if you want to know which host acts as a mail exchanger (MX) for a given domain, you can query a database for all MX resource record entries. The most common resource records are A (host) resource records and PTR (pointer) resource records. The objective of an A resource record is to associate a particular host name with an IP address, whereas a PTR resource record associates an IP address with a particular host name, a process known as reverse lookup.

> **Note** The DNS servers used by Active Directory must support SRV (service) resource records, which are explained later in greater detail.

Zones

Ideally, the first DNS server should maintain the information for the entire tree. However, because the content size of the Internet namespace can result in more records than a single DNS server can maintain, it may be necessary to delegate the authority for a specific portion of a namespace to another DNS server. Delegation in this case is the process of adding an entry that points to another DNS server into the database on the first DNS server. The second DNS server has the authority for the records representing the contiguous set of DNS names that follow the pointer. The contiguous set of records a DNS server is responsible for is called a *zone.* A DNS server that contains a zone is said to be authoritative for the names in that zone. For more information about zones, see the *Windows* 2000 *Sewer Deployment Planning Guide.* For more information about DNS, see http://www.microsoft.com.

Active Directory Logical Structure

Active Directory is an enterprise-wide directory service that is used by both Windows 2000 and Exchange 2000. For organizations that use earlier versions of Exchange, there may be design considerations for Active Directory that are specific to Exchange. It is important to understand these design decisions when you are planning the namespace.

When you plan for and deploy Active Directory in an enterprise, you are defining a significant part of your company's network infrastructure. The way in which you design Active Directory determines:

- The availability and fault tolerance of Active Directory.
- The network usage characteristics of Active Directory clients and servers.
- How efficiently you can manage the contents of Active Directory.
- The way users view and interact with Active Directory.
- The ability of Active Directory to evolve as your Exchange organization evolves.

Design considerations affect each part of Active Directory. These considerations can be broken into four major areas: the Active Directory domain structure, the forest, the organizational unit structure, and Windows 2000 site topology.

Domain Design

Each forest that you create contains one or more domains. Defining the domains that make up your forest, including the domains that support your enterprise, requires:

- Determining the number of domains in each forest
- Choosing a forest root domain
- Understanding the impact of changes to the domain plan after deployment

How you design your domain also determines the availability of Active Directory on the network, the traffic characteristics of client queries, and the traffic characteristics of domain controller replication.

Determining the Number of Domains in Each Forest

To determine the number of domains you need in each forest, first consider having a single domain, even if you currently have more than one domain in your Windows NT 4.0 installation. Next, provide a detailed justification for each additional domain. Every domain that you create introduces some incremental cost in terms of additional management overhead. For this reason, be certain that each domain you add to a forest serves a beneficial purpose.

When to Create More Than One Domain

Three possible reasons for creating additional domains are:

- Preserving existing Windows NT domains

  If you already have domains running Windows NT, you might prefer to keep them as they are instead of consolidating them into a smaller number of domains using Active Directory. Before you decide to keep a domain, be sure to weigh the costs against the long-term benefits of having fewer domains.

- Administrative partitioning

  More domains might be necessary to meet the administrative and policy requirements of your company. Consider the following issues:

  - Unique security policy requirements

    You might want a set of users on your network to abide by a domain user security policy that is different from the security policy applied to the rest of the user community. For example, you might want your administrators to have a stronger password policy, such as a shorter password change interval, than the regular users on your network. To do this, you must place your administrators in a separate domain.

- Autonomous domain administration requirements

  The members of the domain administrators group in a given domain have complete control over all objects in that domain. If you have a division in your company that will not allow outside administrators to have control over their objects, place those objects in a separate domain. For example, for legal reasons, it might not be advisable for a subdivision of a company that works on highly sensitive projects to accept domain supervision from a higher-level Information Technology group. Remember that all domains in the forest must share the Configuration container and schema.

- Physical partitioning

  Physical partitioning involves taking the domains you have in a forest and dividing them up into a greater number of smaller domains. Having a greater number of smaller domains allows you to optimize replication traffic by replicating objects only to places where they are most relevant. For example, in a forest containing a single domain, every object in the forest is replicated to every domain controller in the forest. This might lead to objects being replicated to places where they are rarely used, which is an inefficient use of bandwidth. For example, a user who always logs on at a headquarters location does not need the user account replicated to a branch office location. You can avoid replication traffic by creating a separate domain for the headquarters location and not replicating that domain to the branch office.

Incremental Costs for an Additional Domain

Each domain in the forest introduces some amount of management overhead. When debating whether or not to add a domain to your domain plan, weigh the following costs against the benefits:

- More domain administrators

  Because domain administrators have full control over a domain, the membership of the domain administrators group for a domain must be closely monitored. Each added domain in a forest incurs this management overhead.

- More domain controller hardware

  In Windows 2000, a domain controller can host only a single domain. Each new domain that you create requires at least one computer, and in most cases requires two computers to meet reliability and availability requirements. Because all Windows 2000 domain controllers can accept and originate changes, you must physically guard them more carefully than you did Windows NT 4.0 backup domain controllers, which were read-only computers. Note that the administration delegation within Active Directory domains reduces the requirement for resource domains. Some remote locations that currently must host two domain controllers (a master user domain and a local resource domain) now require only one domain controller, if you choose to consolidate to fewer Active Directory domains.

- More trust links

  For a domain controller in one domain to authenticate a user from another domain, it must be able to contact a domain controller within the second domain. This communication represents an added possible point of failure if, for example, the network between the two domain controllers is malfunctioning at the time. The more users and resources located in a single domain, the less an individual domain controller must rely on being able to communicate with other domain controllers to maintain service.

- Greater chance of having to move users and groups between domains

  The more domains you have, the greater the chance you have to move users and groups between two domains. For example, a business reorganization or a job change for a user can create the need to move a user between domains. To users and administrators, moving a user or group between organizational units inside a domain is a trivial and transparent operation.

  However, moving users and groups between domains is more involved and can impact the user.

- Group Policy and access control do not flow between domains

  Group Policy and access control applied within a domain do not flow automatically into other domains. If you have policies or delegated administration through access control that is uniform across many domains, they must be applied separately to each domain.

Choosing a Forest Root Domain

After you have determined how many domains you will place in your forest, you need to decide which domain will be the forest root domain. The ***forest root*** domain is the first domain that you create in a forest. The two forest-wide groups, Enterprise Admins and Schema Admins, reside in this domain.

Note If all of the domain controllers for the forest root domain are lost in a catastrophic event, and one or more domain controllers cannot be restored from backup, the Enterprise Admins and Schema Admins groups will be permanently lost. There is no way to reinstall the forest root domain of a forest. If your forest contains only one domain, that domain is the forest root domain. If your forest contains two or more domains, consider the following two approaches for selecting the forest root domain:

- Using an existing domain

  From the list of domains you have, select a domain that is critical to the operation of your company and make it the forest root domain. Because you cannot afford to lose this domain, it already requires the kind of fault tolerance and recoverability that is required for a forest root domain.

- Using a dedicated domain

  Creating an additional, dedicated domain to serve solely as the forest root domain carries all the costs of an extra domain, but it has certain benefits that might apply to your company, such as:

  - The domain administrator in the forest root domain can manipulate the membership of the Enterprise Admins and Schema Admins groups. You might have administrators who require domain administrator privileges for some part of their duties, but you do not want them to manipulate the forest-wide administrators groups. By creating a separate domain, you avoid having to place these administrators into the domain administrators group of the forest root domain.
  - Because the domain is small, it can be easily replicated anywhere on your network to provide protection against geographically centered catastrophes.
  - Because the only role the domain has is to serve as the forest root domain, it never risks becoming obsolete. On the other hand, when you select a domain from your planned list of domains to be the forest root domain, there is always a chance that the domain you choose will become obsolete, perhaps due to a change in your company. However, you will never be able to fully retire such a domain, because it must serve as the forest root domain.

Changing the Domain Plan After Deployment

Domain hierarchies are not easy to restructure after they have been created. For this reason, it is best not to create domains that are based on a temporary or short-lived organizational structure. For example, creating a domain that maps to a particular business unit in your company might create work for you if that business unit is split up, disbanded, or merged with another unit during a corporate reorganization.

However, there are cases where organization-based partitioning is appropriate. Geographic boundaries provide a relatively stable template for partitioning, but only if the organization does not frequently move across those boundaries. Consider a domain plan for an army, where the army has different divisions spread across a number of bases. It might be common for divisions to move between bases. If the forest were partitioned according to geographic location, administrators would have to move large numbers of user accounts between domains when a division moved between bases. If the forest were partitioned according to divisions, administrators would only have to move domain controllers between bases. In this case, organization-based partitioning is more appropriate than geographic partitioning.

Exchange 2000 and Active Directory Domain Design

The Windows 2000 Active Directory domain structure has more impact on Exchange 2000 than Windows NT 4.0 had on Exchange version 5.5. In Windows 2000 Server, the domain boundaries define the namespace and each domain includes one or more domain controllers.

It is possible for a Windows 2000 domain controller to be placed into a Windows NT 4.0 domain; this is called a mixed-mode domain. It is important to determine whether the environment is running in mixed mode or in native mode. In a mixed-mode environment, universal groups cannot be used, which impacts the use of security and distribution groups.

In addition, it is important to consider the domain tree structure and verify that Exchange fits properly in it and complies with security and administrative requirements. For example, consider a company that has four separate domains. (Technically, this company could have achieved the same results by using a single domain model and organizational units; however, the decision to create separate domains depends on political or geographical factors.) Multiple domains affect the Exchange design in a number of ways. One impact is user management and migration. Moving users between domains is not as transparent to the user as would be the case if the user were to be moved between servers within a domain. Users migrating from Lotus Notes to Exchange also need more consideration when migration is done in a multiple domain environment.

DNS Service

Clients in an Active Directory environment rely on the availability of the DNS service. Clients query the DNS service for the IP addresses of service providers for the Lightweight Directory Access Protocol (LDAP) and the Kerberos V5 services. Clients require at least one DNS server to locate the domain controllers that support elementary services.

In addition, clients need to locate at least one domain controller for the logon procedure. The client sends a query for a list of the domain controllers in its site to the DNS server. The client contacts all of the domain controllers that the DNS server sends back, and uses the first domain controller that responds to the client's logon procedure.

Advantages of DNS Service

The Windows 2000 DNS service provides the best integration with Active Directory. Additional benefits of using the Windows 2000 DNS service are:

- Increased fault tolerance
- Easier management
- Easier security administration
- More efficient replication of large zones

Windows 2000 Server and Berkeley Internet Name Domain (BIND) 8 implementation support dynamic update. Service (SRV) resource records are supported by the following DNS servers:

- Windows 2000 Server
- Windows NT Server 4.0 Service Pack 4 and later
- BIND 4.9.6 and later

Security

Security is an elementary requirement of the highest importance for a modern computer network. Security guards against all types of illegal network access, including malicious attacks or user mistakes and interference. Security of the objects in a zone database is vital because a DNS database is a complete information system that can provide an intruder with the information needed to attack your network.

Dynamic Update

The Windows 2000 DNS service supports the concept of secure dynamic updates. Dynamic update, which is specified in RFC 2136, relieves administrators of the responsibility of keeping the entries in the zone databases consistent. If you are using the Windows 2000 DNS service, dynamic update allows an administrator to specify which users and computers are allowed to perform changes on the DNS zone databases. In addition, entries in a zone database can have access control lists (ACLs) assigned to them. However, if you are not using the Windows 2000 DNS service, dynamic update is not supported and zone database entries cannot have ACLs assigned to them.

> **Note** Clients do not support dynamic update and must be configured to keep their records up to date. When a Dynamic Host Configuration Protocol (DHCP) server is integrated with dynamic update, it is possible to specify which part of the DHCP server is covered by dynamic update. The administrator can also specify the strategy to be used in the event a name conflict occurs. Unless specified otherwise, clients solve name conflicts by overriding an entry in a zone database.

Using Non-Windows 2000 DNS Service

Active Directory can operate without a Windows 2000 DNS server. However, using a Windows 2000 DNS server makes administration and maintenance easier. If you already have an implemented DNS structure that is running on a third-party DNS server, consider whether the advantages described earlier are worth the expenditure of a dedicated Windows 2000 DNS server. Also, compare the features supported by your DNS server with those features offered by a Windows 2000 DNS server.

If you do not use the Windows 2000 DNS service, you must plan the roles of the available DNS servers very carefully. The main DNS database that contains the records for a given company is known as the ***primary zone***. An administrator can create secondary zones on other DNS servers where the content of the zones is determined by a backup mechanism known as "zone transfer." Using zone transfer guarantees fault tolerance.

> **Note** RFC 9095 introduces the concept of "Incremental Zone Transfer" which reduces network traffic and minimizes the amount of data that must be transferred over the network to keep DNS server content consistent. Incremental zone transfer enables DNS servers to request only those changes that are necessary to keep the zone database up to date. However, the incremental zone transfer environment needs to be well planned.

If you use the Windows 2000 DNS service, it is not necessary to use primary and secondary zones and incremental zone transfer. Instead of configuring the backup roles of the DNS servers, it is possible to connect all DNS servers to a given Active Directory structure, perform zone database changes on each of them, and let the Active Directory replication mechanism update the content of the zone databases. Thus, by using the Windows 2000 DNS service, you can reduce administrative costs dramatically and maintain a very high level of fault tolerance.

Although it is possible to use Active Directory without Windows 2000 DNS servers, you must delegate the following zones to DNS servers if Windows 2000 DNS servers are not used:

- _tcp.*Active* Directory domain name
- _udpActive Directory domain name
- _msdcs.*Active* Directory domain *name*

For more information about Active Directory and DNS, see the Microsoft Windows 2000 Server Resource Kit *TCP/IP* Core Networking Guide, available from Microsoft Press.

Placing DNS Servers

At a minimum, place at least one DNS server in each Windows 2000 Server site. Depending on other considerations, such as the network bandwidth and resultant response times, it may be necessary to add additional DNS servers to your company.

Important The DNS service can run on a domain controller.

In order to reduce replication traffic, DNS servers can be configured as caching-only servers. A caching-only server forwards queries to DNS servers and builds up its cache from the resolved queries of which it is notified. A caching-only server is not responsible for a specific zone.

Besides the amount of network traffic, security must be considered. For example, it might be necessary to place a DNS server in front of a firewall if you don't want your main name server to be accessible from outside your company.

Domain Naming Recommendations

To create the domain hierarchy in a forest, assign a DNS name to the first domain, and then decide for every subsequent domain if it is a child of the first domain or if it is a new root domain. Based on that evaluation, assign names accordingly. Some recommendations for naming domains include:

- Use names relative to a registered Internet DNS name.

  Names registered on the Internet are globally unique. If you have one or more registered Internet names, use those names as suffixes for the domain names in Active Directory.

- Use Internet standard characters.

  Internet standard characters for DNS host names are defined in RFC 1123 as 'A'-'Z', 'a'-'z', '0'-'9', and '-'. Using only Internet standard characters ensures that Active Directory complies with standards-based software. To support the upgrade of domains in Windows NT 4.0 or earlier that have non-standard names, Microsoft client computers and the Windows 2000 DNS service support almost any Unicode characters in a name.

- Never use the same name twice.

  Never give the same name to two different domains, even if those domains are on unconnected networks with different DNS namespaces. For example, if Northwind Traders decides to name a domain on the intranet nwtraders.tld, it should not also create a domain on the Internet called nwtraders.tld. If a nwtraders.tld client computer connects to both the intranet and Internet simultaneously, it selects the domain that answers first during the SRV record search. To the client, this selection appears random, and there is no guarantee that the client selects the intended domain. An example of such a configuration is a client computer that has established a virtual private network connection to an intranet over the Internet.

- Use names that are distinct.

  Some proxy client software, such as the proxy client built into Microsoft Internet Explorer or the Windows Sockets proxy client, uses the name of a host to dctcnnine if that host is on the Internet. Most software of this type provides, at minimum, a way of excluding names with certain suffixes as being local names, instead of assuming they are on the Internet.

  If Northwind Traders wants to call an Active Directory domain on their intranet nwtraders.tld, they must enter nwtraders.tld in the exclusion list of their proxy client software. This prevents clients on the Northwind Traders intranet from seeing a host on the Internet called www.nwtraders.tld.

To avoid this problem, Northwind Traders can use a registered name that does not have a presence on the Internet, such as nwtraders-internal.tld, or establish a company policy that requires names ending in a specific suffix or nwtraders.tld. Thus, for example, corp.nwtraders.tld would never appear on the Internet. In both cases, it is easy to configure proxy client exclusion lists so that they can determine which names are on the intranet and which are on the Internet.

There are many different techniques for accessing the Internet from a private intranet. Before using any name, ensure that it can be properly resolved by client computers on your intranet given your specific Internet access strategy.

- Use the fewest number of trees possible.

  There are some advantages to minimizing the number of trees in your forest. The following advantages might apply in your environment:

  - After you have been given control over a particular DNS name, you own all names that are subordinate to that name. The smaller the number of trees, the smaller the number of DNS names that you must own in your company.
  - There are fewer names to enter in the proxy client exclusion list.
  - Non-Microsoft LDAP clients might not use the global catalog server when searching Active Directory. Instead, to perform directory-wide searches, these client computers use ***deep*** *searches.* A deep search covers all of the objects in a particular subtree. The fewer trees in a forest, the fewer deep searches that are required to search the entire forest.

- Make the first part of the DNS name the same as the NetBIOS name.

  It is possible to assign entirely unrelated DNS and NetBIOS names to a domain. For example, the DNS name of a domain could be sales.nwtraders.tld, but the NetBIOS name could be "Marketing."

  Computers that are not running Windows 2000 Server and Windows 2000 software will display and accept NetBIOS names, whereas computers running Windows 2000 Server and Windows 2000 software will display and accept DNS names. This can lead to confusion on the part of your users and administrators. You should only use unmatched NetBIOS and DNS names if:

  - You want to migrate to a new naming convention on your network.
  - You are upgrading a NetBIOS name that contains non-standard characters but you want the DNS name to have all standard characters.

- Review names internationally.

  Names that have a benign or useful meaning in one language can sometimes be derogatory or offensive in another language. DNS is a global namespace; be sure to review your company's names globally.

  If you have multiple localized versions of Windows running on your network, all computers, including Windows 2000 Professional and all versions of Windows 2000 Server, must use only Internet-standard characters in both their DNS and NetBIOS names. If you use characters other than Internet-standard characters, only computers with the same locale setting can communicate with each other.

- Use names that are easy to remember.

  Typically, administrators see domain names by using Windows 2000 Server administrative tools. Users typically work with global catalog servers that can be queried without knowledge of a domain or host name. However, users sometimes use the domain name. For example, domain names are used for user principal name logons or if the domain namespace matches an SMTP address. In general, choose domain names with components that are easy to remember.

- Differentiate domain names and computer names.

  In Windows NT 4.0 and earlier, a computer is identified primarily by a NetBIOS name, which is the name by which the computer is known on the network. In Windows 2000, a computer is identified primarily by its full computer name, which is a DNS fully qualified domain name (FQDN). The same computer can be identified by more than one FQDN. However, only the FQDN that is a concatenation of the host name and the primary DNS suffix is the full computer name. By default, the primary DNS suffix of a computer that is running Windows 2000 Server is set to the DNS name of the Active Directory domain to which the computer belongs. The primary DNS suffix can also be specified by Group Policy. For more information about Group Policy, see the Windows 2000 Server documentation.

DNS and Exchange 2000

Although Windows 2000 domains use a DNS-style naming convention, a domain name such as europe.microsoft.tld does not dictate the SMTP address for Exchange mailbox-enabled users created within that domain. Although a user's logon name might be someone@europe.microsoft.tld, the e–mail address generation is controlled by recipient policies in the organization. Unlike Exchange 5.*x*, Exchange 2000 can automatically generate multiple addresses for users.

You can configure multiple recipient policies for the organization. Each one has an LDAP filter rule based on RFC 2254 associated with it. This allows administrators to generate proxy addresses, such as SMTP, based on fields within Active Directory. For example, all research and development personnel in the company can have a different (or additional) SMTP address from other members of the company.

Where possible, align the user's domain logon name with the SMTP address to reduce potential user confusion. You might need to use a combination of recipient policies and a user principal name to achieve this, as seen in Table 4.1.

Table 4.1 Domain logon name and SMTP address alignment

Property	Value
Residing domain	example.microsoft.tld
Recipient policy	@microsoft.tld
SMTP address	firstname.lastname@microsoft.tld
Pre-Windows 2000	example\firstname.lastname
User principal name logon	firstname.lastname@microsoft.tld

Note The Recipient Update Service generates SMTP and proxy addresses. The Recipient Update Service checks whether any other object in the forest has the same address, based on mail nickname and SMTP domain/proxy root. If a duplicate is found, the Recipient Update Service appends a number to the duplicate address to make it unique. For example: bob@microsoft.tld, bob2@microsoft.tld, and bob3@microsoft.tld. When users are created, Active Directory verifies that the user principal name is unique. Active Directory does not allow you to create a user if the user principal name is not unique.

Configuring a Unified Namespace

Recipient policies are tightly integrated with DNS. To configure a unified namespace, first configure DNS to identify the appropriate Exchange 2000 server by creating a mail exchanger (MX) record for each e-mail domain you plan to define in your organization. For example, if you have three Exchange servers and each one processes incoming mail for multiple departments, and each user has multiple valid SMTP addresses, you need to define each Exchange server in DNS with a mail exchanger (MX) record to identify the three domains to be handled by a particular Exchange server. Table 4.2 shows what your DNS records might look like for the zone nwtraders.microsoft.

Table 4.2 DNS record example

Department	Record Type	Record Value	Location
London	A	172.19.240.2	
Tokyo	A	172.19.241.2	
Seattle	A	172.19.242.2	
Engineering	MX	10	London
Personnel	MX	10	London
Sales	MX	10	London
Manufacturing	MX	10	Seattle
Headquarters	MX	10	Seattle
Exporting	MX	10	Tokyo
Support	MX	10	Tokyo

In this scenario, all mail sent to user@nwtraders.microsoft.tld will be delivered to the London server. Backup mail deliveries will be made to Tokyo and Seattle, based on the priority. Additionally, mail sent to user@engineering.nwtraders.microsoft.tld, user@personnel.nwtraders.microsoft.tld, and user@sales.nwtraders.microsoft.tld will be delivered to the London server. Mail sent to .tld and user@headquarters.nwtraders.microsoft.tld will be delivered to Seattle. Mail sent to user@exporting.nwtraders.microsoft.tld and user@support.nwtraders.microsoft.tld will be sent to Tokyo.

User Principal Name

A user principal name is an e-mail-like name that uniquely represents a user. A user principal name consists of two parts, a user identification portion and a domain portion. The two parts are separated by an "@" symbol, to form *user@DNS-domain-name*, for example, suzan@nwtraders.microsoft.tld. Every user is automatically assigned a default user principal name, where the user portion of the name is the same as the user's logon name, and the *DNS-domain-name* portion of the name is the DNS name of the Active Directory domain where the user account is located. When logging on using a user principal name, users no longer have to choose a domain from a list on the logon dialog box.

You can set user principal names to arbitrary values. For example, even if Suzan's account is in the nwtraders.microsoft.tld domain, her user principal name could be set to suzan@nwtraders.tld. When the user logs on, the user account to be validated is discovered by searching the global catalog for a user account with a matching user principal name value. By making user principal name values independent from domain names, administrators can move user accounts between domains, leaving user principal name values unchanged and making moves between domains more transparent to users.

The user principal name logon is an attribute of the Active Directory account and can be set up so that a user can log on to the network by a short and recognizable name, thus hiding the complexity of the underlying domain infrastructure. The user principal name attribute is replicated to the global catalog.

Some companies choose to make their user's logon alias and SMTP alias the same so that users need to remember only one alias when logging on or when referring to their SMTP address. Other companies purposefully avoid having the user logon name and alias be the same and make it as difficult as possible for hackers by making each user's SMTP alias and logon alias different.

Organizational Unit Structure

Objects in Active Directory exist within domains. They can be further subdivided within a domain by using organizational units. While the domain defines the replication and security context, the organizational units within the domain define the location of objects and how they are administered, and provide a method of applying Group Policy to different groups of objects in an efficient way.

Organizational Unit Characteristics

The following characteristics of organizational units are important to consider when creating structure in a domain:

- Organizational units can be nested.

  An organizational unit can contain child organizational units, enabling you to create a hierarchical tree structure inside a domain.

- Organizational units can be used to delegate administration and to control access to Active Directory objects.

  When you use a combination of organizational unit nesting and access control lists, you can delegate the administration of objects in the directory in a very detailed manner. For example, you could grant a group of help desk technicians the right to reset passwords for a specific set of users, but not the right to create users or modify any other attribute of a user object.

- Organizational units are not security principals.

  You cannot make organizational units members of security groups, nor can you grant users permission to a resource because they reside in a particular organizational unit. Because organizational units are used for delegation of administration, the parent organizational unit of a user object indicates who manages the user object, but it does not indicate the resources a user can access.

- Group Policy can be associated with an organizational unit.

  Group Policy enables you to define desktop configurations for users and computers. You can associate Group Policy with sites, domains, and organizational units. Defining Group Policy on an organizational unit basis allows you to use different policies within the same domain.

- Users will not navigate the organizational unit structure.

  It is not necessary to design an organizational unit structure that will appeal to users. Although it is possible for users to navigate the organizational unit structure of a domain, it is not the most efficient way for a user to discover resources. The most efficient way to find resources in the Active Directory is by querying the global catalog.

Organizational Unit Planning Process

Create an organizational unit structure plan to document the specific reason for creating each organizational unit for a domain. Note the specific reason for creating an organizational unit each time you add one to the plan. This will help you make sure that every organizational unit has a purpose, and it will help the readers of your plan to understand the reasoning behind the structure.

The steps to creating an organizational unit structure for a domain are:

1. Create organizational units to delegate administration.
2. Create organizational units to hide objects.
3. Create organizational units for Group Policy.
4. Understand the impact of changing organizational unit structures after deployment.

It is important to create organizational units in the order presented above. An organizational unit structure designed specifically for delegating administration is shaped differently from an organizational unit structure designed specifically for Group Policy. Because there are multiple ways of applying Group Policy, but only one way to delegate administration, create organizational units for delegating administration first.

Creating Organizational Units to Delegate Administration

In Windows NT, delegation of administration within a domain was limited to the use of built-in local groups, such as the account administrators group. These groups had predefined capabilities, which in some cases did not fit the needs of a particular situation. As a result, there were situations where administrators in a company needed high levels of administrative access, such as domain administrator rights.

In Windows 2000 Server, delegation of administration is more powerful and flexible. This flexibility is achieved through a combination of organizational units, per-attribute access control, and access control inheritance. Administration can be delegated arbitrarily by granting a set of users the ability to create specific classes of objects or modify specific attributes on specific classes of objects.

Delegating administration in your company has several benefits. Delegating specific rights enables you to minimize the number of users who must have high levels of access. Accidents or mistakes made by an administrator with restricted capability will only have an impact within the administrator's area of responsibility. Previously, in your organization it might have been necessary for groups other than IT to submit change requests to high-level administrators, who would make these changes on their behalf. By delegating administration, you can delegate responsibility to the individual groups in your organization and eliminate the overhead of sending requests to high-level administrative groups.

Three ways to delegate administration are:

- By physical location. For example, administration for objects in Europe can be handled by an autonomous set of administrators in Europe.
- By business unit. For example, administration of objects belonging to the engineering department can be handled by an autonomous set of administrators in the engineering department.
- By role or task. For example, a set of administrators might be responsible for computer account objects.

Modifying Access Control Lists

The access control entries (ACEs) in the access control list (ACL) of an object determine who can access that object and what kind of access they have. When an object is created in the directory, a default ACL is applied to it. The default ACL is described in the schema definition of the object class. To delegate administration, grant a group specific rights over an organizational unit by modifying the ACL of the organizational unit.

ACEs can be inherited by child objects of a container object. If any of the child objects are also containers, the ACEs are applied to the children of those containers as well. With inheritance, you can apply a delegated right to an entire subtree of organizational units instead of a single organizational unit. You can also block ACE inheritance on an object to prevent ACEs from a parent container from applying to that object or any child objects. Inheritable ACEs apply only within a domain and do not flow down to child domains.

Always reference groups in ACLs, not individual users. Managing the membership of a group is simpler than managing an ACL on an organizational unit. When users change roles, it is much easier to discover and change their group memberships than to check the ACLs on every organizational unit. Where possible, delegate to local groups instead of global or universal groups. Unlike global groups, local groups can have members from any trusted domain, making them better suited for granting resource permissions. Unlike universal groups, local group membership is not replicated to the global catalog, making local groups less resource intensive.

Creating Organizational Units to Hide Objects

Even if a user does not have the right to read the attributes of an object, that user can still see that the object exists by enumerating the contents of that object's parent container. The easiest and most efficient way to hide an object or set of objects is to create an organizational unit for those objects and limit the set of users who have the List Contents right for that organizational unit.

Creating Organizational Units for Group Policy

In Windows NT 4.0, you can use the System Policy Editor to define user and computer configurations for all of the users and computers in a domain. With Windows 2000 Server, you use Group Policy to define user and computer configurations, and associate those policies with sites, domains, or organizational units. Whether or not you need to create additional organizational units to support the application of Group Policy depends on the policies you create and the administrative delegation options you select.

Changing the Organizational Unit Structure After Deployment

Moving an object or subtree of objects changes the parent container of those objects. ACEs that were inherited from the old parent no longer apply, and there might be new inherited ACEs from the new parent. To avoid unexpected changes in access, evaluate in advance what the changes will be and determine whether those changes will have any impact on the users that currently access and manage those objects.

Moving a user object, computer object, or a subtree containing user or computer objects can change the Group Policy that is applied to those objects. To avoid unexpected changes in client configurations, evaluate the changes in Group Policy and ensure that they are acceptable for users.

Exchange 2000 and Organizational Units

The location of a user or server within a domain does not affect Exchange 2000. This means that regardless of how you choose to design your organizational unit structure, whether you base it on administrative delegation requirements or Group Policy requirements, Exchange 2000 is not affected.

Active Directory objects that can be mailbox-enabled do not have to be in the same organizational unit or even in the same domain as the Exchange 2000 server on which the mailbox physically resides. With Exchange 2000 and Active Directory, moving an object from one organizational unit to another is a simple operation that does not affect the associated mailbox. Mailboxes can be moved between mailbox stores on an Exchange server and between servers, independently of whether the user object is moved in Active Directory.

In some companies running earlier versions of Exchange, Exchange servers are members of Windows NT resource domains. During the migration from Windows NT to Windows 2000 Server it is likely that resource domains will be collapsed into Active Directory–based domains that contain both users and resources. These companies may choose to locate the objects in Active Directory in a single domain organizational unit. However, this will have little effect on the Exchange server or how it is managed. Administration of server objects can be delegated separately from administrator permissions on the server itself.

Note The organizational unit structure is not shown in the Exchange 2000 address book.

Forest

A forest is a collection of one or more Windows 2000 Active Directory trees, organized as peers and connected by two-way transitive trust relationships between the root domains of each tree. All trees in a forest share a common schema, configuration, and global catalog. When a forest contains multiple trees, the trees do not form a contiguous namespace.

Important The forest represents the boundary for the Exchange 2000 organization. An Exchange 2000 organization cannot span multiple Windows 2000 forests. If a company has multiple forests, multiple Exchange organizations are required.

When multiple domains exist in the forest, you can change the mode of each domain one at a time. Therefore, you can switch your parent or child domains to native mode without the other being in native mode.

When a collection of domains and domain trees are joined together to form a single Active Directory, this is called a forest. A forest uses a single schema and configuration definition, which is replicated to all domain controllers in every domain.

Number of Forests

The decision to use more than one forest has significant resource and functionality effects. Carefully consider the best course of action for your situation.

Creating a Single-Forest Environment

A single forest environment is simple to create and maintain. All users see a single directory through the global catalog, and do not need to be aware of the underlying directory structure. When adding a new domain to the forest, no additional trust configuration is required. Configuration changes need to be applied only once to affect all domains.

Creating a Multiple-Forest Environment

Because forests have shared elements, such as schemas, it is necessary for all the participants in a forest to agree on the content and administration of those shared elements. Organizations such as partnerships and conglomerates might not have a central body that can drive this process. In short-lived organizations like joint ventures, it might not be realistic to expect administrators from each organization to collaborate on forest administration.

If administration of your network is distributed among many autonomous divisions, it might be necessary to create more than one forest. If your company chooses to create multiple forests, it is important to understand the implications of having multiple forests and the limitations this imposes on users and objects.

Although non-transitive trust relationships can be implemented between domains that reside in different forests (including Windows NT 4.0 domains), a multiple-forest environment can present certain challenges to Exchange 2000 designs.

Exchange 2000 and the Forest

The Active Directory forest determines the boundaries of the Exchange 2000 organization. It is not possible to have Exchange 2000 servers within the same Exchange 2000 organization in different forests.

There are two ways to implement Exchange 2000:

- In a single forest by using Active Directory, knowing that all domains trust one another.
- In multiple forests by using Active Directory, establishing coexistence between earlier versions of Exchange and Active Directory.

In a multiple-forest environment, an Exchange 2000 organization must be created for each forest. This has the following effects on Exchange 2000:

- There is more than one Exchange 2000 organization to administer.
- There is no automatic Active Directory replication between multiple forests. Therefore each forest has a separate global address list. (Users can see only one global address list).
- You cannot configure routing group connectors between Exchange 2000 organizations. (You must use SMTP connectors or X.400 connectors instead).
- No link state information transfers between Exchange 2000 organizations because routing group connectors cannot connect organizations.

Deploying Active Directory Forests

Having multiple forests causes additional administrative work and can affect the deployment of Exchange 2000. In this scenario, you can create manual trusts between domains in separate forests. However, these domains are non-transitive, which means that you can have a domain model that resembles a Windows NT 4.0 deployment, with multiple manual trusts between domains. Additionally, the global catalog servers only show objects within their own forest; this determines what Microsoft Outlook users see when they perform searches.

If you deployed Exchange in the past, you could create a single Exchange Server 5.5 organization across two Windows NT 4.0 domains with no trusts between them. This organization could include multiple sites and be deployed across these two domains without relying on the underlying security structure. Because Exchange Server 5.5 performed mail-based directory replication between sites, users in both domains were able to see all users within the same organization through one global address list.

Synchronizing Data Between Forests

If Outlook users and Exchange 2000 servers are deployed in multiple forests, you may need to replicate certain data between the two systems to achieve the functionality you need.

You can use Microsoft Metadirectory Services or third-party products, such as the Compaq LDAP Synchronization Utility (LDSU), ISOCOR's MetaConnect, or Siemens' DirX. For more information about Microsoft Metadirectory Services, see "Inter-Organization Replication and Directory Synchronization" in this book and the Microsoft Web site at http://www.rnicrosoft.com.

> **Important** Synchronization of directory data between forests changes the target object class. For example, a mailbox-enabled user in organization A is a mail-enabled contact in organization B.

Exchange 2000 includes the Public Folder Inter-organization Replication Tool, which can synchronize public folders between different organizations. However, replicating public folders between different organizations creates additional administration overhead compared to replicating public folders in a single organization.

The Public Folder Inter-organization Replication Tool can also replicate information such as Free/Busy System folders. This information allows users from different forests to schedule meetings with one another and view free and busy time slots on their calendars.

> **Important** There is no automatic solution for replicating users' calendars between organizations; therefore, users cannot open calendars that exist in another organization.

Windows 2000 Site Topology

A Windows 2000 site is a local, logical collection of Internet Protocol (IP) subnets. All computers that are in the same site have high-speed connectivity—local area network (LAN) speeds—with one another. Unlike an Exchange 5.x site, a Windows 2000 site does not correspond to a specific part of the namespace. For example, multiple sites can exist within a single domain, and conversely, a single site can span multiple domains. Synchronous remote procedure calls (RPC) replicate the domain naming context within a domain.

Windows 2000 sites help to define the physical structure of a network. When a change occurs in Active Directory, sites can be used to optimize replication traffic and to enable users to connect to a domain controller by using a reliable, high-speed connection.

Distinguishing Windows 2000 Sites from Exchange 5.5 Sites

Windows 2000 sites only define zones in the underlying network where high bandwidth is present. A site in Exchange 5.5 or earlier defines the unit of administration and namespace and dictates message routing. Unlike the Exchange 5.5 site, a Windows 2000 Server site does not correspond to a specific part of the namespace and is not a partition for administration.

Active Directory Logical and Physical Structure

In Active Directory, the logical structure is separate from the physical structure. You use the logical structure to organize your network resources, and you use the physical structure to configure and manage your network traffic. The physical structure of Active Directory is composed of sites and domain controllers and defines where and when replication and logon traffic occur. Understanding the physical components of Active Directory is critical to optimizing network traffic and the logon process. In addition, this information can help in troubleshooting replication and logon problems.

Domains define the logical structure of your company, whereas sites define the physical structure of your network. The logical and physical structures of Active Directory are independent of each other. This means:

- There is no necessary correlation between your network's domain structure and its physical structure.
- Active Directory allows multiple domains in a single site and multiple sites in a single domain.
- There is no necessary correlation between domain namespaces and sites.

Creating a Windows 2000 Site Topology

A Windows 2000 site topology describes a physical network for a forest. Creating the site topology requires taking the physical topology of your network and describing it in terms of available bandwidth and network reliability. Active Directory clients and servers use the site topology of a forest to route queries and replication traffic efficiently. A site topology also helps you decide where to place domain controllers on your network.

When you create your Windows 2000 site topology, it is useful to have a complete map of the physical topology of your network. That map should include the list of physical subnets on your network, the media type and speed of each network, and the connections between each network.

Site and Site Topology Information

Sites, site links, and subnets are all stored in the configuration container, which is replicated to every domain controller in the forest. Every domain controller in the forest has complete knowledge of the site topology. A change to the site topology causes replication to every domain controller in the forest.

> **Note** Site topology is separate and unrelated to domain hierarchy. A site can contain many domains, and a domain can appear in many sites.

Keep the following key concepts in mind when designing your site topology:

- Create a site for each LAN, or set of LANs, that is connected by a high-speed backbone, and assign the site a name. Connectivity within the site should be reliable.
- Create a site for each location that does not have direct connectivity to the rest of your network and is only reachable by SMTP mail.
- Determine which sites will not have local domain controllers, and merge those sites with other nearby sites. Sites help efficiently route traffic between clients and domain controllers, and between domain controllers and other domain controllers. Without a domain controller in a site, there is no traffic to be routed.
- Client computers attempt to communicate with domain controllers in the same site as the client before trying to communicate with domain controllers in any other site. Anytime bandwidth between a set of networks is plentiful enough that you do not care if a client on one network communicates with a server on a different network, then consider those networks all to be in one site.
- If your entire network consists of fast, reliable connectivity, the entire network can be considered to be a single site.
- Anytime two networks are separated by links that are heavily used during parts of the day and idle during other parts of the day, those networks should be put into separate sites. You can schedule replications between sites to prevent replication traffic from competing with other traffic during high usage hours.

  Note If a client computer is on a subnet that is not defined in Active Directory. it is not considered part of a site, and it selects randomly from all domain controllers for a given domain. You may encounter situations where not all subnets are defined in Active Directory, such as when new subnets are being added to your network.

Site Links

Site links are used to model the amount of available bandwidth between two sites. As a general rule, any two networks connected by a link that is slower than LAN speed are considered to be connected by a site link. A fast link that is near capacity has a low effective bandwidth and can also be considered a site link. Site links have four parameters:

- **Cost** The cost value of a site link helps the replication system determine when to use the link when compared to other links. Cost values determine the paths that replication takes through your network.
- **Replication schedule** A site link has an associated schedule that indicates at what times of day the link is available to carry replication traffic.
- **Replication interval** The replication interval indicates how often the system polls domain controllers on the other side of the site link for replication changes.
- **Transport** The transport that is used for replication.

Connect sites with site links to reflect the physical connectivity of your network. Assign each site link a name. Site links are transitive, so if site A is connected to site B, and site B is connected to site C, then the Knowledge Consistency Checker will assume that domain controllers in site A can communicate with domain controllers in site C. You need to create a site link between site A and site C only if there is, in fact, a distinct network connection between those two sites. A backbone network that connects many sites can be represented by a single site link that connects many sites, instead of by separate links between sites. Reducing the number of site links that need to be created and managed is particularly useful when many sites have the same characteristics.

For each site link, record the following:

- Link speed and current usage levels.
- Whether the link is pay-by-usage.
- Whether the link is historically unreliable.
- Whether the link is only intermittently available.

Exchange 2000 and Windows 2000 Site Design

Whereas Windows 2000 sites group domain controllers together for purposes of replication and client access, Exchange 2000 employs routing groups to group Exchange 2000 servers together so that message routing can be managed and scheduled.

> **important** Routing groups perform a function similar to Exchange sites in earlier versions of Exchange. Be careful to distinguish between Windows 2000 sites and Exchange sites in earlier versions of Exchange; they serve different purposes.

Message Routing and Group Expansion

All message routing information, including routing groups and bridgeheads, is held within the configuration naming context of Active Directory. To determine routing, the Exchange 2000 server contacts a local domain controller and retrieves this information.

If a message is sent to a universal group, the SMTP virtual server, which is configured to perform the expansion, uses LDAP to contact a global catalog and populate the message header with the group membership. If the message is for a domain local or global group, the expansion server should be in the same domain as that group and should be configured to use only global catalogs from the local domain where the group resides. By default, every Exchange server tries to use global catalogs from its local domain and site; however, if there are not enough global catalogs in the local domain and site, Exchange requests global catalogs from any domain in the local site. To ensure correct expansion of domain local and global groups, any expansion server for these types of groups should be located in a Windows 2000 site that contains only global catalogs from the same domain as both the expansion server and the groups.

Placement of Sites and Routing Groups

Whether you are defining Windows 2000 sites or Exchange routing groups, you must:

- Have permanent and reliable network connections between all servers.
- Provide adequate network bandwidth between servers.
- Focus on reducing communication latency within the site and network bandwidth utilization between sites.

Note The location of domain controllers in your site topology has a direct effect on the availability of services provided by Active Directory.

Trust Relationships

Active Directory provides security across multiple domains through trust relationships between domains. When there are trust relationships between domains, the authentication mechanism for each domain trusts the authentication mechanism for all trusted domains. If a user or application is authenticated by one domain, its authentication is accepted by all other domains that trust the authenticating domain. Users in a trusted domain have access to resources in the trusting domain, subject to the access controls that are applied in the trusting domain.

Note Access to resources in any discussion of trust relationships always assumes the limitations of access control. Trust relationships allow users and computers to be authenticated by an authentication authority. Access control allows authenticated users to use the resources that they are authorized to use and prohibits them from using resources that they are not authorized to use.

Transitive and Non-transitive Trust

In Windows 2000 Server, domains can be joined to a domain tree or forest, and each child domain has an automatic two-way trust relationship with the parent domain. This trust relationship is also transitive. Transitive trust means that the trust relationship extended to one domain is extended automatically to any other domain that is trusted by that domain. Transitive trust is applied automatically for all domains that are members of the domain tree or forest. Therefore, when a grandchild domain is created, the trust relationship between the parent and child domains is accepted by the grandchildren domains, and vice versa. For example, if a user account is authenticated by the parent domain, the user has access to resources in the grandchild domain. Similarly, if the user is authenticated by a child domain, the user has access to resources in the parent domain, as well as in the grandparent domain.

The effect of transitive trust in Windows 2000 Server is that there is complete trust between all domains in an Active Directory forest—every domain has a transitive trust relationship with its parent domain, and every tree root domain has a transitive trust relationship with the forest root domain.

Complete trust makes managing multiple domains simpler in Windows 2000. In previous versions of Windows NT, a popular model for deploying domains was the multiple master domain model. In that model, a domain containing primarily user accounts was called a master user domain, and a domain that contained primarily computer accounts and resources was called a resource domain. Adding a new domain to the deployment required several trusts to be created. With Active Directory, when you add a domain to a forest it is automatically configured with two-way transitive trust. This eliminates the need to create additional trusts with domains in the same forest.

Limiting Trust Relationships

If there are two entities in a forest within a conglomerate or partnership that do not trust one another, there is no way of breaking the trust between domains. Also, there are accounts and groups within the forest that have rights to affect the forest. These groups, such as Exchange Administrators, contain users that are trusted across all partners or companies. It might be necessary to create more than one forest if:

- Individual entities do not trust each other's administrators.

  Entities such as partnerships and conglomerates may not be able to decide on administrative roles and responsibilities. Or they may not be able to agree on the content and administration of shared elements, such as the schema. In short-lived entities like joint ventures, it might not be realistic to expect administrators from each entity to collaborate on forest administration.

- Individual entities cannot agree on a forest change policy

  Schema changes, configuration changes, and the addition of new domains to a forest have forest-wide impact. Each of the participating entities in a forest must agree on a process for implementing these changes, and on the membership of the schema administrators and enterprise administrators groups. If participating entities cannot agree on a common policy, they cannot share the same forest.

- You want to limit the scope of a trust relationship.

  Every domain in a forest trusts every other domain in the forest. Every user in the forest can be included in a group membership or appear on an access control list on any computer in the forest. If you want to prevent certain users from ever being granted permissions to certain resources, then those users must reside in a different forest from the resources. If necessary, you can use explicit trust relationships to allow those users to be granted access to resources in specific domains.

Multiple domains within Active Directory are linked to a domain tree. Users in the linked domains can access resources in other domains by way of transitive trust relationships that exist hierarchically among all of the domains in the domain tree. However, administrative rights are not inherently transitive. Thus, an additional layer of security is added by limiting the scope of domain administrative rights.

Optimizing Authentication with Shortcut Trust Relationships

When a user requests access to a network resource, a domain controller from the user's domain must communicate with a domain controller from the resource's domain. If the two domains are not in a parent-child relationship, the user's domain controller must also communicate with a domain controller from each domain in the trust tree between the user's domain and the resource's domain. Depending on the network location of the domain controllers for each domain, each extra authentication hop between the two domains can increase the chance of a possible failure or increase the likelihood of authentication traffic having to cross a slow link. To reduce the amount of communication necessary for such interactions, you can connect any two domains with a shortcut trust relationship.

For example, if you have multiple trees in a forest, you might want to connect the group of tree roots in a complete mesh of trust. Remember that in the default arrangement, all tree roots are considered children of the forest root from a trust perspective. That means authentication traffic between any two domains in different trees must pass through the forest root. Creating a complete mesh of trust allows any two tree root domains to communicate with each other directly.

Active Directory Logical Components

Active Directory consists of components that function to organize it into a logical structure. This section describes three of these logical components: naming contexts, the global catalog server, and groups.

Naming Contexts

A naming context is a self-contained section of the Active Directory hierarchy that can have its own properties, such as replication configuration and permissions structure. Active Directory uses naming contexts to define the boundaries for information held within the database structure. The information stored in Active Directory on every domain controller in the forest is partitioned into three categories: domain, configuration, and schema data. These naming contexts are the units of replication in Active Directory. If the domain controller is also a global catalog server, it holds a fourth category of data as well. In multi-domain forests, domain controllers belonging to different domains have a common configuration and schema naming context, but they have a domain unique domain naming context.

Domain Naming Context

The domain naming context contains all of the objects in Active Directory for a domain. Domain data in each domain is replicated to every domain controller in that domain, but not beyond that domain. Domain objects include recipient objects, such as users, contacts, and groups.

Configuration Naming Context

The Configuration container is replicated to every domain controller in the forest. For example, Active Directory stores information about the physical network in the Configuration container and uses it to guide the creation of replication connections between domain controllers. The enterprise administrators security group has full control over the Configuration container. Sharing a single, consistent configuration across the domains of a forest eliminates the need to configure domains separately.

The configuration container contains replication topology and related metadata. Applications that use Active Directory, like Exchange 2000 Server, store the configuration of the Exchange 2000 organization in the Configuration container. Because Active Directory replicates the configuration container between all domains in the forest, the configuration of the Exchange 2000 organization is also replicated throughout the forest. The Configuration container defines the topology, connectors, protocols, and service settings for the Exchange 2000 organization.

Schema Naming Context

The schema naming context defines the object classes and the attributes of object classes that can be created in Active Directory. Object classes are the types of objects that can be created in Active Directory. The schema naming context contains all object types (and their attributes) that can be created in Active Directory. This data is common to all domains in the forest, and is replicated by Active Directory to all domain controllers in the forest. The schema administrators security group has full control over the schema.

Exchange 2000 and Naming Contexts

During the installation of the first Exchange 2000 server in the forest, the Active Directory schema is extended with new attributes for Exchange 2000, which have names that start with ms-Exch- attribute — for example, ms-Exch-OAB-Default. Simultaneously, existing Active Directory attributes are modified, some of which affect how Outlook presents information. Many LDAP Data Interchange Format (LDIF) files are imported as part of the installation process for the first Exchange 2000 server in Active Directory. This process can take an extended period of time to complete.

Active Directory is used to store Exchange 2000 data such as recipient objects, configuration data, schema attributes, and the global address list. A separate directory for Exchange is no longer necessary because Exchange 2000 is fully integrated with Active Directory. Exchange 2000 stores the majority of its information in the configuration container naming context. The configuration container naming context is replicated to every domain controller in the forest, simplifying global administration.

> **Important** Because Exchange 2000 uses Active Directory, it is critical to review and understand the Windows 2000 naming conventions. Any issues with the naming convention should be resolved with Windows administrators. Attributes that are relevant to only Exchange, such as e–mail address attributes, should be defined and added to the naming convention.

Global Catalog Server

A global catalog is a Windows 2000 Server domain controller that stores a writable copy of the domain naming context for its domain and the forest-wide configuration and schema naming contexts. The global catalog also contains a select set of the attributes for every object from every domain in the forest. In addition, it provides a complete replica of the configuration and schema naming contexts for the forest.

By default, the partial set of attributes stored in the global catalog includes those attributes most frequently used in search operations, because one of the primary functions of the global catalog is to support clients querying Active Directory. The global catalog makes directory structures within a forest transparent to users and enables fast, efficient searches that span the entire forest.

The availability of global catalog servers is crucial to the operation of Active Directory. For example, a global catalog server must be available when processing a user logon request for a native-mode domain, or when a user logs on with a user principal name.

When processing a logon request for a user in a native-mode domain, a domain controller sends a query to a global catalog server to determine the user's universal group memberships. Since groups can be explicitly denied access to a resource, complete knowledge of a user's group memberships are necessary to enforce access control correctly. If a domain controller of a native-mode domain cannot contact a global catalog server when a user wants to log on, the domain controller refuses the logon request.

Global Catalog Server Placement

It is very important to understand how Active Directory usage affects global catalog server load, and how to efficiently deploy your global catalog servers. Use the same fail-over and load distribution rules that you used for individual domain controllers to determine whether additional global catalog servers are necessary in each site.

For best results, monitor the use of the global catalog by Exchange 2000 and add servers as necessary to gain optimal performance. Because LDAP is a standard protocol and the LDAP data is not encrypted, it is possible to view the traffic going to and from an Exchange 2000 server and the global catalog server by using Network Monitor (unless Exchange 2000 is installed on the global catalog server).

Exchange 2000 balances requests between available global catalog servers. Each time an address book search occurs and each time a message is routed, Exchange uses the closest global catalog server. When deciding on the number of global catalog servers required to support an Exchange 2000 deployment, consider the power of your servers, the number of users, the volume of messages sent, and other factors that affect your users and the load carried by your servers.

> **Note** In a single domain environment. global catalog servers are not required to process a user logon request. However, as a general rule, you should designate at least one domain controller in each site as a global catalog server. Client computers still seek global catalog servers for search operations. Also, having global catalog servers already in place allows your system to adapt gracefully if you add more domains later.

Exchange 2000 and Global Catalog Server

Users running Outlook 2000 directly query a global catalog server to get addressing information. For other clients, Exchange 2000 acts as a proxy to the global catalog server. Global catalog servers are the primary locations for users to access information about Active Directory and Active Directory services available for Exchange 2000.

It is important to review which fields are set to be replicated to the global catalog servers and update that list based on your organization's requirements. For example, the Department attribute is not replicated by default. However, there may be a need to make the Department attribute available due to a workflow application that depends on that attribute.

Exchange 2000 uses LDAP to query and update Active Directory, making heavy demands on the global catalog servers. For this reason, global catalog servers should be strategically positioned so that Exchange 2000 can perform quick searches and users can access resources with minimum delays. Exchange 2000 has a directory access mechanism that allows it to share the results of LDAP queries among the Exchange 2000 components.

> **Important** Exchange 2000 servers should be installed as member servers of the domain and not as domain controllers or global catalog servers, mainly for performance reasons. A server that is a domain controller or global catalog uses a lot of processing power and memory, competing with Exchange 2000 for the same processing power and memory. However, if you have sufficient hardware, you can use an Exchange 2000 server as a domain controller or global catalog, especially if the Exchange server is lightly used or if there is a need in remote office scenarios to consolidate servers onto fewer, more powerful hardware systems.

Clients can access Active Directory information by communicating directly with a global catalog server or by using DSProxy. Active Directory supports both Messaging Application Programming Interface (MAPI) and LDAP queries (for backward compatibility with older MAPI clients). Exchange 2000 provides a communication process for Microsoft Outlook Web Access clients making directory queries using HTTP. For more information about DSProxy, see "Active Directory Integration and Replication" in this book.

Global Address List

When Outlook users want to find a person within the organization, they usually search the global address list (GAL), which represents an aggregation of all messaging recipients in the enterprise. Because Exchange 2000 servers no longer host their own directory service, all data is retrieved from the global catalog servers in Active Directory. Because a global catalog server can support the MAPI protocol as well as LDAP, Outlook clients can communicate with Active Directory using the same protocol employed by the Exchange Server 5.5 directory service.

Getting Directory Data

Depending on the type of request being made, an Exchange *2000* server can go to different Active Directory servers to retrieve data. An Exchange *2000* server usually establishes a number of LDAP connections to nearby domains controllers and global catalog servers. An Exchange *2000* server performs two main Active Directory activities: address book searches and configuration data searches. For an address book search, the Exchange *2000* server queries the closest available global catalog server.

When an Exchange *2000* server needs to read configuration data such as routing information, it can connect to any domain controller within the local domain to get this information. This is possible because the configuration container naming context is replicated to every domain controller in the forest. An Exchange *2000* server would rather use the same domain controller for these requests than switch between different domain controllers.

Active Directory Groups

In Exchange *2000*, an Active Directory group is the equivalent of an Exchange Server *5.x* distribution list. Active Directory contains two types of groups: a security group and a distribution group. Each group type has a scope attribute. The scope of a group determines who can be a member of the group, and where you can use that group in the network. Three group scopes are available: domain local, global, and universal. The domain mode limits the choice of group type and group scope.

Group Types

Both security groups and distribution groups can be mail-enabled and used as the equivalent of distribution lists. Use these group types to create distribution lists for the various groups within your company. In addition, consider the following factors when working with groups:

- **Group Names** The naming standards for mail-enabled groups should follow the Windows *2000* naming standards.
- **Group Ownership** Group ownership is assigned to the user that requested the group in order to reduce administrative overhead.
- **Group Size** Keep the number of members in a group of either type below *5,000*. Although this is not a hard limit within Active Directory, a listing of *5,000* should work in all circumstances. For efficiency and scalability, you should consider nesting groups that are above *500* members.

Security Groups

Security groups provide the following advantages:

- They can be mail-enabled by assigning an SMTP address, allowing the group to act as a distribution list equivalent.
- They can be used for assigning permissions to public folders in Exchange 2000.
- They can be useful if you want to also assign network permissions to the group members.
- They can lower the number of groups and the amount of maintenance required for Active Directory and the messaging system, because they can act as pseudo-distribution groups.

Security groups provide the following disadvantages:

- Users might gain unauthorized access to network resources if they are accidentally placed in the wrong group.

Distribution Groups

Distribution groups provide the following advantages:

- They can be used for bulk mailing.
- They can be used as universal groups even in a mixed-mode domain.

Distribution groups provide the following disadvantages:

- Permissions cannot be assigned to network resources.
- Permissions cannot be assigned to public folders in Exchange 2000.

Group Scopes

Universal groups offer the greatest flexibility in a company. However, because their membership is stored in the global catalog server, changes to universal groups are replicated between all global catalog servers in the company. For this reason, it is recommended that universal groups remain fairly static.

Domain local or global groups are the best choice for groups with dynamic membership because group replication occurs only within the domain. However, because their membership is not included in the global catalog server, these groups are not visible to users in other domains.

Domain Local Scope

Groups with the domain local scope have the following attributes:

- In a native-mode domain, they can contain user accounts, global groups, and universal groups from any domain in the forest, as well as domain local groups from the same domain.
- In a mixed-mode domain, they can contain user accounts and global groups from any domain.
- You can grant permissions to domain local groups only for objects within the domain in which the domain local group exists. Permissions cannot be assigned to network resources or public folders in other domains.
- They can be converted to a universal group when they exist in a native-mode domain, as long as they do not contain another domain local group.
- The group object is listed in the global catalog, but the group membership is not.
- Outlook users in other domains cannot view the full membership.
- Group membership must be retrieved on demand if expansion takes place in a remote domain.

Global Scope

Groups with the global scope have the following attributes:

- They can contain user accounts from the same domain and global groups from the same domain, when in native-mode domains.
- They can contain user accounts from the same domain, when in a mixed-mode domain.
- You can grant permissions to global groups for all domains in the forest, regardless of the location of the global group.
- They can be converted to a universal group when in a native-mode domain, as long as they are not a member of any other global group.
- They can only contain recipient objects from the same domain.
- The group object is listed in the global catalog, but the group membership is not.
- Outlook users in other domains cannot view the full membership.
- Group membership must be retrieved on demand if expansion takes place in a remote domain.

Universal Scope

Groups with the universal scope have the following attributes:

- They can contain user accounts from any domain, global groups from any domain, and universal groups from any domain in the forest, when the domain is in native mode.
- Universal security groups can only be used in native-mode domains; universal distribution groups can be used in mixed-mode and native-mode domains.
- You can grant permissions to universal groups for all domains in the forest, regardless of the location of the universal group.
- They cannot be converted to any other group scope.
- Outlook users in any domain can view full membership.
- Membership never has to be retrieved from remote domain controllers.
- Membership modifications are replicated to the global catalog servers.

Limiting the membership of universal groups to groups only, rather than individual user accounts, enables you to adjust the user accounts that are members of the universal group by adjusting the membership of the groups that are part of the universal group. Because this does not directly affect the membership of the universal group, no replication traffic is generated.

Mail-Enabling Groups

Suggested guidelines for mail-enabling a group are:

- Use security groups as the primary type, but only mail-enable the groups when appropriate.
- Use distribution groups for lists that include non-trusted recipients.
- Use universal groups when the ability to view membership is important. In Windows 2000 mixed mode, distribution groups must be used.

Each type of distribution group can be further divided into three categories, depending on whether you want the group to be accessible to specific domains or to all users. Exchange 2000 servers can route mail to group members regardless of which group or category is used. The group type used is more relevant when considering how Outlook users will view them and how access to secure resources will be granted. For more information on recipients and groups, see the Windows 2000 Server and Exchange 2000 Server documentation.

Exchange 2000 and Groups

The type and scope of the group that you use for Exchange 2000 depends on your business and user requirements. For full flexibility, implement universal security groups. Although this is a security group definition, it can be mail-enabled by adding an SMTP address and viewed in the global address list. However, a disadvantage of universal security groups is that they can only be created in native-mode domains. You can move from mixed-mode to native-mode domain by upgrading the domain controllers to Windows 2000 Server. Moving to native-mode domains simplifies the upgrade and deployment process for Exchange 2000 and provides additional directory scalability.

Another point to consider when deploying universal groups is that their membership is listed in the global catalog servers, so any membership change causes replication traffic. Although Active Directory supports property-level replication, the membership for a group is held in a multi-valued property on the group object. Therefore, if the group is large, a significant amount of replication traffic can be generated. To mitigate this risk, place user objects in other universal groups and then nest these groups under an umbrella universal group. When the membership changes for a user in the group, the large universal group object is not changed and no replication traffic is created.

When you implement universal groups, Outlook users can still view full membership of both the umbrella group and its subgroups. You can use global groups instead of universal groups. However, global groups do not have their membership listed in the global catalog servers and this can impact the ability of an Outlook client to view membership at the recipient level.

When a message is sent to a group, the SMTP service must expand the membership of the group object. If it is a domain local or global group defined in the local domain, the membership list can be retrieved from any local domain controller. In addition, if it is a universal group and users appear directly on the list, the membership can be obtained from any local global catalog server.

If a message is sent to a domain local or global group that has been created in another domain, or if a universal group contains global groups that are in other domains, then the group must specify an expansion server that exists in its domain. That expansion server must use the global catalog from the same domain as the local or global group in order for the expansion to be successful.

A decision has to be made as to whether it is better to create a universal group or use domain local and global groups and set the membership retrieval to remote domains when the group needs to be expanded. Consider the following questions when deciding upon the group scope:

- Should there be single or multiple domains in Active Directory?
  - For single domains, a universal group is not needed because all domain objects are local.
  - For multiple domains, use universal scope for fairly static groups. However, keep in mind that users might not have access to all object attributes from other domains in universal groups.
- Is direct IP connectivity possible between all domains?
  - If the answer is yes, and if the group is static and wants to include users from other domains, use universal scope.
  - If the answer is no, use local domain or global scope, especially if the group is dynamic. Users need complete access to all attributes of group members.
- Does membership change often?
  - If the answer is yes, use local domain or global scope.
  - If the answer is no, use universal scope.
- Does the majority of the e-mail to the group come from local or remote domains?
  - If the answer is local domains, use local domain or global scope.
  - If the answer is remote domains, use universal scope.

Keep in mind that, when you implement domain local or global groups, Outlook users cannot view the membership of the group unless you define the domain local or global groups within the same domain as the user.

You must also remember that groups are used for determining public folder access. Unlike previous versions of Exchange, the store does not need to expand a group when a user accesses a public folder. Because all ACLs for public folders are based in Active Directory, group membership is carried in a user's access token and is presented to the Exchange server upon connection to the resource.

> **Note** When the administrator attempts to create a new group in a mixed-mode domain, by default it is configured as a security group with a domain local scope. A new group in a native-mode domain defaults to a security group with a universal scope. In mixed-mode domains, you cannot change a group's scope after you have created it.

Group Creation Decision Process

Implementing the correct group type for Exchange 2000 mailing and public folder access is very important. Before you create the group, make sure you fully understand how it will be used. The scope, size, and membership of the group are key factors when implementing different group types. Use the scenarios in the next section to help guide you through your decision process.

Coexisting with and Upgrading from Earlier Versions of Exchange

In a coexistence scenario, the Active Directory Connector (ADC) replicates Exchange distribution lists to Active Directory groups. In Exchange 2000, Exchange Server 5.x distribution lists replicate to Active Directory as a universal distribution group.

If the ADC has not been configured, the Exchange 2000 server upgrade process converts any existing distribution lists in Exchange 5.5 to universal distribution groups if the Active Directory domain is in mixed mode, and universal security groups if the Active Directory domain is in native mode.

Exchange 2000 Scenarios

Requirement To allow the sales team in Toronto to e–mail one another with lead information.

Solution Create a global distribution.

Reasoning If all of these members are in the same domain, you can use a global group. The purpose of the group is to enable members to send e–mail to each other, so access to network resources is not a consideration. If access to network resources is a consideration, create a security group instead. This will prevent having two groups with the same membership if requirements change in the future. The list will be used only by team members in the same domain, so membership does not need to be retrieved from remote domain controllers.

Requirement The worldwide product marketing unit wants to have a list created so that anyone within the company can e–mail them. Users do not need to know who is on the mailing list.

Solution Create a domain local distribution group.

Reasoning You cannot use global groups because the members of the group are in a multitude of domains. Because network access is not desired and the membership does not need to be published to the global catalog, you do not need to create a universal group. The marketing unit also sees frequent changes to their group; therefore, it is not appropriate to create a universal group. However, to prevent membership from being remotely retrieved when a user in a remote domain sends a message to the group, the list is set to expand on a server in the domain in which the group is defined.

Requirement To create a small mail-based discussion forum for a new company product and to give the team access to all of the Web sites and network shares that contain information about the product.

Solution Create a universal security group.

Reasoning Members of the discussion forum can be located anywhere in the world. Because there is a strong possibility that members are in different domains, it is necessary to use universal scope for the group. The team is quite small and it is not envisaged that membership will change on a frequent basis, so it is appropriate to put the user membership directly within the universal group. This increases usability because members can see who is on the list before they send sensitive or confidential information. The group created is a security group because network resource access is also a requirement. Note that the domain in which the universal group is created must be in native mode.

Requirement To create a mailing list to mail security announcements and important information in bulk to the entire company.

Solution Create a global distribution group in every user domain, and then nest these global groups into a universal group.

Reasoning Access to network resources is usually controlled on a per-team basis or through the default access control entry, which grants permissions for users who have not been explicitly defined on the ACL. As a result, it is extremely unlikely that you will need to grant permissions on network resources for these groups. Membership for this group will change fairly frequently, making it inappropriate to populate full membership into a single universal group. When a membership change takes place, only the domain controllers within that domain need to be updated with the modification. Global catalog servers do not need to replicate any data. From a user's viewpoint, it is extremely unlikely that they will want to see the membership of this group. The global groups are nested within the universal group, so it is very easy for users to e–mail the entire company (permissions permitting) by using a single address instead of selecting each individual team or region. Another advantage of this method is that users can send e-mail to an individual group very easily. (Note that domain local groups cannot be used in this scenario because they cannot exist within universal groups.)

CHAPTER 5

Active Directory Integration and Replication

Patrick McFarland, Senior Consultant, Microsoft
Tony Soper, Product Manager, Microsoft
Walden Barcus, User Education Manager, Microsoft

With the advent of Active Directory directory service in Microsoft Windows 2000, Microsoft has integrated the Exchange directory database with the Windows 2000 Active Directory so that the Windows 2000 operating system now contains all user account details, security information, and messaging data. This eliminates a need for a separate directory for Microsoft Exchange. Because of this unification, Exchange 2000 is now considered an Active Directory integrated application.

In earlier versions, Exchange used its own database to hold messaging system and directory information, including mailbox, custom recipient, distribution list, and public folder directory objects. Although mailbox objects were linked to the Microsoft Windows NT security database through the primary Windows NT account, no other link between the operating system and messaging directory existed. The primary reason for this connection was that the Windows NT Security Accounts Manager was not designed to hold rich information attributes, such as telephone numbers and certificates.

This chapter explores the unification of the Exchange directory and Active Directory and how this unification affects planning for Exchange directory services, replication, and coexistence.

In This Chapter

Active Directory Replication
Global Catalog Servers
Active Directory Services
Client Access to Active Directory
Coexistence and Upgrading

Active Directory Replication

In Active Directory, the directory tree represents all the objects in a Windows 2000 forest, and is partitioned in a way that allows it to be distributed to domain controllers in different domains within a forest. The Active Directory replication model encompasses how changes are propagated and tracked among domain controllers. Each domain controller in a forest stores copies of particular parts of the directory, and each defined segment of the directory is a *directory partition.* A copy of the contents, of one directory partition, on a domain controller is called a *replica.* Updates to replicas are synchronized among the domain controllers that store the same directory partitions during replication.

A partial replica of the information within each domain replicates to global catalog servers between domains. This read-only, partial replica contains a subset of the attributes of all directory partition objects.

Active Directory uses *multi-master* replication to synchronize directory information. Multi-master replication enables multiple replicas to exist, all of which can be independently updated by applications or administrators at any time. Those changes automatically propagate to all the replicas, but if the system reaches a steady state with no replica updates, then they will eventually converge to the same state. This multi-master replication is similar to the directory synchronization of Exchange 5.5, but unlike other directory services that use a master server and subordinate server approach to providing updates, where all updates must be made to the directory master copy and then replicated to the subordinate copies. With Active Directory, no specific domain controller is the master server. Instead, all domain controllers within a domain are equivalent. Changes can be made to any domain controller and then replicated to other domain controllers, according to the Active Directory replication topology.

Multi-master updating provides highly available access to write to directory objects because several servers contain updateable copies of an object. Each domain controller in the domain accepts updates independently, without communicating with other domain controllers. The system resolves any conflicts in updates to a specific directory object. If updates cease and replication continues, all copies of an object eventually contain the same value.

> **Note** Do not use Active Directory as if it were an Exchange directory. This does not make optimal use of Active Directory. Table 5.1, from the Windows 2000 Deployment Guide, examines the differences between Active Directory replication and those of Exchange 5.5 Directory Service.

Table 5.1 Active Directory and Exchange 5.5 Directory Service comparison

Active Directory	Exchange Server 5.5 Directory Service
Master replicas accept object updates independently.	Each directory service object is mastered in a specific site. Multi-master updates are possible within the master site.
The basis for replication is the object GUID. When an object is renamed, its GUID does not change, so renaming the object cannot lead to replication errors.	The basis for replication is distinguished names. Therefore, to avoid problems, Exchange does not rename objects.
Replicates an update by transmitting only the changed attributes.	Replicates an update by transmitting the entire object.
Supports replication data compression between sites over remote procedure call (RPC) or Simple Mail Transfer Protocol (SMTP) transports.	Supports compression of replication data between sites over SMTP transport only.
Supports servers that contain only a subset of the objects in the entire directory. Has global catalog servers that contain all objects, but only a partial set of attributes.	Maintains a complete replica of the directory on each directory server. The schema in Exchange is site-specific and not replicated out of its site.
Has a flexible replication topology (including choice of transports).	Has a replication topology between sites, limited to a tree structure that cannot contain redundant links. Replication transport between sites is limited to e-mail.
Uses Windows 2000 sites to help generate replication topology and to help clients perform intelligent replica selection; however, sites are not tied to directory partitioning.	Uses Exchange sites to generate replication topology. Sites are also the unit of directory partitioning.

For general information about the directory replication of Windows 2000, see the *Microsoft Windows* 2000 *Sewer Resource Kit Deployment Planning Guide* available from Microsoft Press.

Global Catalog Servers

A global catalog is a specific type of Windows 2000 domain controller. All domain controllers expose port 389, and contain an updateable copy of the directory and configuration partitions of the domain. If the server is the schema master, it also contains an updateable copy of the schema partition.

A global catalog also exposes port 3268, with objects from every partition, but only a subset of the object attributes, including objects in the global catalog's own domain. Thus, even though an object is available on ports 389 and 3286, only a partial set of attributes is visible on 3286. The entire object can be seen on port 389 of the domain controller for the domain that contains the object. MAPI properties appear only on global catalogs. Outlook users within a forest can see directory details for users in other domains because MAPI properties replicate to the global catalog.

In the Windows 2000 environment, users access directory and directory service information from global catalog servers. Since these servers are critical to Exchange 2000 Server, this section examines sizing requirements, server placement, and replication concerns.

Determining the Number of Global Catalog Servers Required

Exchange 2000 load balances requests among available global catalog servers. When you are deciding on the number of global catalog servers needed to support an Exchange 2000 deployment, you should consider the number of clients supported by each global catalog server.

For scalability and resilience, put into place at least two global catalog servers per Windows 2000 site. If a Windows 2000 site spans multiple domains, configure a global catalog server for each domain where Exchange 2000 servers and clients are situated.

The number of global catalog servers depends on the capabilities of your servers, number of users, volume of messages sent, and other factors that affect processor load. For best results, monitor the demands of Exchange 2000 on the global catalog and then add servers to gain optimal performance. Microsoft provides several tools and methods to help you estimate the number of global catalog servers for your environment. Some of these tools and methods are described here.

Global Catalog Database Sizing

As you determine the number of global catalog servers you need, you should understand the factors that affect the size and performance of the global catalog servers within your Windows 2000 environment. Primary factors critical to determining this are:

- Database size on the domain controllers
- Database size on the global catalog servers
- Bandwidth requirements for replication

Use the Active Directory Sizer tool to find the estimate for the database size on the domain controllers; download the tool from the Microsoft Web site at http://www.microsoft.com.

To estimate database size for global catalog servers, use the following table. The total size is an estimate for the initial default database.

Table 5.2 Object sizes In the Initial global catalog sewer database

Object Type	Number of Objects	KB per Object	Total for Class In KB
Active Directory initial database	I	12,000	12,000
User with mandatory attributes	Users	4.366	Total KB = Number of users: 4.366 KB
Additional attribute	Additional attributes: users with additional attributes	0.100	Total KB = Number of attributes: users with additional attributes: 0.100 KB
Contact	Contacts	1.678	Total KB = contacts: 1.678 KB
Group	Groups in the forest (mail enabled or security)	2.097	Total KB = groups: 2.097 KB
Group member in small group	Members in small groups (20 members)	0.200	Total KB = members in small groups: 0.200
Group member in medium group	Members in medium groups (100 members)	0.100	Total KB = members in medium groups: 0.100
Group member in large group	Members in large groups (more than 200 members)	0.070	Total KB = members in large groups: 0.070
Organizational unit	Organizational units in forest	1.992	Total KB = organizational units: 1.992 KB
Certificate	Certificates in Active Directory	2.181	Total KB = certificates: 2.181 KB
Computer	Computer accounts in the forest	4.086	Total KB = computers: 4.086 KB
Volume	Volumes registered in Active Directory	1.573	Total KB = volumes: 1.573 KB
Printer	Printers in the forest	2.412	Total KB = printers: 2.412 KB
BLOB[1]	BLOBs	File size + 25 bytes	Total KB = Total of BLOB file sizes + (Number of BLOBs: 0.025)

'A BLOB (binary large object) – a large file, typically image or sound, that must be handled in a special way because of its size.

After you calculate the kilobytes needed (see Table 5.2) you will have a good representation of the total database space required for the global catalog servers from all domains in your forest. The local domain controller database containing default global catalog data (without additional attributes and schema extensions) has approximately 60 percent of the total kilobytes calculated, representing the global catalog information for that single domain. Take 60 percent of the total space to calculate the total global catalog overhead needed on each domain controller in each domain. Keep in mind that the default global catalog database size does not take into account the fragmentation of the database or tombstone objects (objects that are deactivated and marked for deletion). You should double the size of your final global catalog database to account for these values. The final database size provides a good estimate of the global catalog servers' database size.

Global Catalog Server Sizing

Estimates for global catalog server sizing suggest that a global catalog can support a fairly large number of users, dependent on their activities and the other available global catalog server services. As a start, provide global catalog services for a remote office of 5 to 25 users within a Windows 2000 site. For a high–end global catalog for a large number of users, server size and number of servers are factors of:

- Number of users
- Performance of the application accessing the global catalog
- Queries per second
- Types of queries (depth, breadth)
- Network availability (latency)
- User expectations

Global Catalog Server Placement

The Windows 2000 recommendation for global catalog servers is to provide a global catalog for each Windows 2000 site. The consideration for Exchange and global catalog server placement is to have an adequate level of service for directory queries. A rough estimate for the relationship between Exchange servers and the global catalog servers (within a data center or concentrated location of Exchange servers) is to have one global catalog server to every four Exchange mailbox servers.

For satisfactory client global catalog access and queries, the number and location for the global catalog servers depends on the same factors as server sizing: number of users, queries per second, types of queries, network availability, and user expectations.

Use Table 5.3 to begin planning the placement of the global catalog servers, given that each server is at least a 500–MHz Pentium III with 512 MB of RAM. However, note that global catalog server placement and sizing must be tested and planned for your environment.

Table 5.3 General rules for planning global catalog servers

Environment	Number of Users	Number of Global Catalogs	Global Catalog Placement
Branch office	1 to 25	1	Branch office, if there are at least 25 users.
Small office	26 to 250	Two for load balancing	Small office
Regional office	251 to 5000	Two for fault tolerance, and additional global catalogs as needed	Regional office
Main office	5000 to 10000	Two for fault tolerance, and additional global catalog servers as needed.	Main office

Selectable Field Replication

With earlier versions of Exchange and Outlook, users rely on the directory data that is populated to the Exchange directory. For example, users can look up each other's telephone numbers through this mechanism.

With Exchange Server 4.0 and Exchange Server 5.0, any directory data present in one Exchange server has to be replicated to every other Exchange server within a company. To reduce network traffic, Exchange Server version 5.5 implements selectable-field replication, which prevents certain attributes (but not objects) from replicating between Exchange sites. Many companies do not take advantage of this new feature because users frequently rely on the directory information.

Active Directory has a similar feature in which selected attributes can be tagged for replication to the global catalog server. A difference between this feature in Active Directory and earlier versions of Exchange is that, by default, not all attributes are tagged for replication. Exchange administrators control attributes tagged for replication through the Active Directory schema console. You can enable the console by typing **Regsvr32 schmmgmt.dll** at the command prompt. Once you enable the console, start the Microsoft Management Console (MMC) process and add the Active Directory Schema snap-in to the console.

To change which attributes replicate to the global catalog, right-click the attribute in the Active Directory Schema, choose **Properties,** and then select or clear the **Replicate this attribute to the Global Catalog** check box, as shown in Figure 5.1.

> **Important** Changing the attributes that replicate affects replication time and network overhead. For more information about replication, see "Making Schema Changes" later in this chapter.

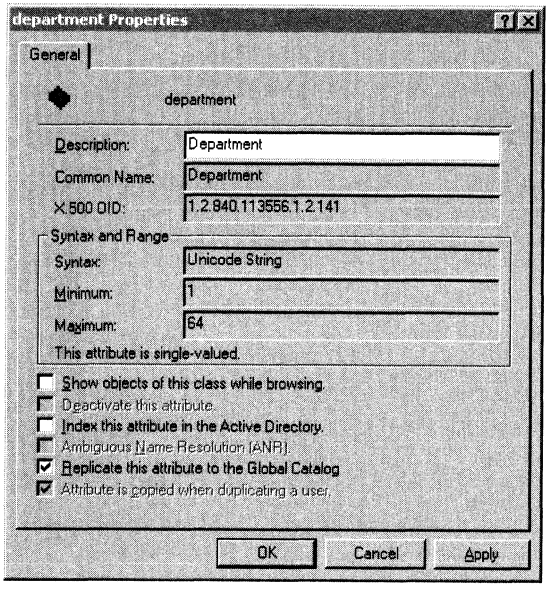


Figure 5.1 Active Directory schema Department attribute properties

Assume you have a forest that contains a root domain and two child domains. Each child domain has a global catalog. All directory attributes have been populated into Active Directory domains, and the *Department* attribute has not been tagged for global catalog replication. Lightweight Directory Access Protocol (LDAP) clients such as Outlook Express, cannot detect the *Department* attribute in either child domain, including their own, because LDAP clients can only see global catalog attributes that are tagged for replication. MAPI clients, such as Outlook 2000, can see attributes when attributes are tagged for replication, and the attributes contain MAPI identifiers.

Selecting Attributes to Replicate to the Global Catalog

Do not remove any functionality that the Outlook users already have if Exchange is currently deployed. Consider the ramifications of replication traffic if you replicate too many additional attributes.

If you have slow network links, you might want to survey Outlook users to find out which directory attributes they use. Find out whether any custom Collaboration Data Objects (CDO) and Active Directory Services Interface (ADSI) applications in your environment require particular Exchange 5.5 attributes. For example, a workflow application might require access to a custom attribute that contains the limit of a manager's budget.

When Microsoft Exchange 2000 Server is installed, it extends the Active Directory schema and tags many of the attributes that users normally require for global catalog replication. However, there are some important attributes, such as *Department*, that are not tagged automatically.

Each additional attribute that you manually tag for replication causes additional replication traffic per object. However, the total replication traffic caused by an existing Exchange *5.x* network is far greater than the replication traffic that Active Directory produces. This is based on the following assumptions:

- Each Exchange *5.x* server in the organization must contain a full copy of the Exchange directory, whereas Active Directory only replicates to domain controllers and global catalog servers.
- Any change to an Exchange *5.x* object causes it to replicate again to the rest of the Exchange organization, whereas Active Directory uses per-property replication, thereby replicating smaller changes.
- The Exchange *5.x* directory replicates around the company network, and Windows NT 4.0 domain controllers also replicate. Active Directory consolidates these two directories into one.

Exchange 2000 Installation and Active Directory Schema

Many LDAP directory interface format files are imported as part of the installation process for the first Exchange 2000 server in Active Directory. This process can take an extended period of time to complete.

When you install Exchange 2000, the Active Directory schema is extended with new attributes that all start with *ms–Exch–*, for example, *ms–Exch–OAB–Default*. Existing Active Directory attributes are modified, some of which affect what Outlook users see in the global address list. Table 5.4 lists the common attributes and LDAP names for an Active Directory, mailbox-enabled, user object. The table illustrates how these attributes differ between standard installations of Active Directory and Active Directory when it is enabled for Exchange 2000. The **In Global Catalog** column indicates whether the attribute is tagged by default for global catalog server replication. The table may be helpful for planning.

Table 5.4 Attribute in Active Directory before and after Exchange 2000 installation

Attribute	LDAP Name	In Index Before	In Global Catalog Before	In Index After	In Global Catalog After
First name	*GivenName*	Yes	No	Yes	Yes
Initials	*Initials*	No	No	No	Yes
Last name	*Sn*	Yes	No	Yes	Yes
Display name	*DisplayName*	Yes	No	Yes	Yes
Alias	*MailNickname*	N/A	N/A	Yes	Yes
Mailing address	*StreetAddress*	No	No	No	Yes
City	*L*	Yes	No	Yes	Yes
State	*St*	No	No	No	Yes
ZIP code	*PostalCode*	No	No	No	Yes
Country or region	*C*	No	No	No	Yes
Job title	*Title*	No	No	No	Yes
Company	*Company*	No	No	No	Yes
Department	*Department*	No	No	No	Yes
Office	*PhysicalDeliveryOfficeName*	Yes	No	Yes	Yes
Telephone	*TelephoneNumber*	No	No	No	Yes
Fax	*FacsimileTelephone Number*	No	No	No	Yes
Home telephone	*HomePhone*	No	No	No	Yes
Manager	*Manager*	No	Yes	No	Yes
SMTP Address	*Mail*	Yes	No	Yes	Yes
Custom attributes (all)	*ExtensionAttribute*-xx	N/A	N/A	No	Yes

Making Schema Changes

The Active Directory schema can grow to accommodate a company's changing needs. Keep in mind that Active Directory schema modifications affect the company's infrastructure and replication mechanisms. For example, if you deploy Active Directory and then decide to tag another attribute for global catalog replication, you can make the change easily in the Active Directory Schema Manager snap-in. However, this causes all global catalogs to set their Update Sequence Numbers to zero. As a result, all objects (not just the changed property) in Active Directory rnust replicate to each global catalog again, consuming significant network bandwidth,

In addition, if you install an application (such as Exchange 2000) that tags attributes for replication in the global catalog, it will have the same impact on Active Directory. Plan your schema to minimize replication time for an Exchange 2000 Server installation.

Without the Exchange 2000 version of Active Directory Connector (ADC), you cannot connect to an Exchange 5.5 environment. For companies that plan to deploy Active Directory before installing Exchange 2000 servers, Exchange-specific schema changes should be imported into Active Directory prior to installation by using the ForestPrep utility **(setup /ForestPrep**) into Exchange 2000. The ForestPrep utility is described in "Exchange Organization Preparation" later in this chapter.

Address Book Searches

In earlier versions of Exchange, each Exchange server holds a copy of the directory. Outlook clients refer to the directory on the server that houses the users' mailboxes, and the message transfer agent (MTA) always uses the local directory to route messages. Thus, in earlier versions of Exchange, each Exchange server corresponds to a directory server, and clients for a given server use only the directory on that server. In Exchange 2000, many Exchange servers can use a single directory server (the global catalog server). With Exchange 2000, address book searches changed because Exchange 2000 no longer has its own directory, and it uses Active Directory for client address book searches. For information about how Exchange 2000 executes address book searches for each messaging client, see "Address Lists and Offline Address Lists" later in this chapter.

Active Directory Services

Because the Exchange 2000 directory is integrated with Active Directory, Exchange 5.5 directory services, such as distribution and address lists, have been removed from Exchange and integrated with Active Directory. As you prepare for Exchange 2000, plan for how these services are delivered through Active Directory. The following services provided by Exchange 2000 and Active Directory are discussed in this section, as well as the coexistence with earlier versions of Exchange:

- Global address list
- Address lists
- Offline address books

Global Address List

Outlook users generally search for other users within their company by means of the global address list (GAL), an aggregation of all messaging recipients. Because Exchange 2000 servers no longer host their own directory services, all data comes from Active Directory through global catalog servers. Because a global catalog server can support the MAPI protocol as well as LDAP, Outlook clients can communicate with Active Directory using the same protocol as the Exchange Server 5.5 directory service.

Exchange Server 5.5 Distribution Lists and Active Directory Groups

Distribution lists in Exchange 5.5 serve two purposes: they send mail to large groups and they assign permissions to public folders. A single object class is used for a distribution list. Because of the directory architecture in earlier versions of Exchange, distribution list membership is replicated to each Exchange server in the organization. By default, all users can view the membership for a given list, although this can be restricted on a per-list basis. It is possible for any Exchange 5.5 MTA to expand the contents of a distribution list, because each is represented on every server.

In Active Directory, a group is the equivalent of a distribution list. There are two types of groups in Active Directory: security and distribution. There are three types of group scope in Active Directory: local, global, and universal. Security groups grant access to network resources, and distribution groups send e-mail to groups. For more information about groups, see "Active Directory Design" in this book.

Address Lists and Offline Address Lists

Address lists are critical to users, especially those who expect to have the functionality that they are accustomed to after they have been migrated to Exchange 2000. The following sections describe differences between earlier versions of Exchange and Exchange 2000.

Address Book Views

Microsoft Exchange Server version 5.0 introduced the concept of address book views. This functionality enables the administrator to create certain views available for Outlook users, based on field groupings. With Exchange Server 5.*x*, when Outlook users search the server-based address lists, they see each site, site container, and the aggregated GAL. For example, to create a virtual container of all users in the Sales Team, the Exchange 5.x administrator creates a view grouped by the field *Department.* Although virtual containers work well, they have various limitations. A major limitation of virtual containers is that when views based on a field are created, one virtual container is created for each unique instance of data within that field. For instance, the creation of a virtual container for Sales Team members would mean that other team containers would automatically be created. This issue exists because the rules that determine view creation are not flexible enough for many large companies.

Address Lists

In Exchange 2000, address book views have been replaced with address lists. You can create a container that has a build rule for the address lists associated with it through the Exchange System Manager console. These rules use the LDAP search filter syntax, as defined in RFC 2254, and they are extremely flexible. For example, if you want to create an address list of all permanent employees in the Toronto Marketing department, you could create a single container called **Marketing** with a rule of:

```
(&(mail=*) (&(department=Marketing)(l=Toronto)(!(Extension-Attribute-
3=Contractor))))
```

Unlike Exchange Server 5.5, where the address list is aggregated with the directory service, the Exchange 2000 GAL is not aggregated within Active Directory because it is also built using a rule.

To ensure that build rules meet your requirements, System Manager allows you to search with the rule when you create it. To verify that the address list is up-to-date, the Recipient Update Service, part of the Exchange System Attendant, polls the directory on a scheduled basis and populates the address lists as necessary. As the user objects are updated, if diagnostics logging is activated, you can find information about the update in event 8174 from the MSExchangeAL service in the application log.

The default address lists, configured in Exchange 2000, appear in the following table.

Table 5.5 Default Exchange 2000 address lists

Address List Name	Syntax Based on RFC 2245
Default global address list	(&(mailnickname=*)(|(&(objectCategory=person)(objectClass=user)(!(homeMDB=*))(!(msExchHomeServerName=*))) (&(objectCategory=person)(objectClass=user)(|(homeMDB=*)(msExchHomeServerName=*)))(&(objectCategory=person)(objectClass=contact))(objectCategory=group)(objectCategory=publicFolder)))
All users	(& (mailnickname=*) (| (&(objectCategory=person)(objectClass=user)(!(homeMDB=*))(!(msExchHomeServerName=*))) (&(objectCategory=person)(objectClass=user)(|(homeMDB=*)(msExchHomeServerName=*)))))
All groups	(& (mailnickname=*) (| (objectCategory=group)))
All contacts	(& (mailnickname=*) (| (&(objectCategory=person)(objectClass=contact))))
Public folders	(& (mailnickname=*) (| (objectCategory=publicFolder)))
All conferencing resources	(msExchResourceGUID=*)

The Recipient Update Service populates the defined address lists by entering the list name and location to the *showInAddressBook* attribute on the user objects in the directory. You can view these values with utilities such as Active Directory Service Interface Editor (ADSIEDIT). In the Active Directory Service Interface Editor, right-click a user and then select **Properties.** On the **Attributes** tab, in **Select which properties to view,** choose **Optional.** In **Select a property to view,** choose **showInAddressBook**.

To install the Active Directory Services Interface Editor, insert the Windows 2000 Server installation CD, open the **\Support\Tools** directory, and run **Setup.** Additional tools are available with the *Microsoft Windows* 2000 *Sewer Resource* Kit, published by Microsoft Press.

Address List Compatibility with Exchange Server 5.x

When you upgrade Exchange Server 5.5 to Exchange 2000, only the default views are created. Customized address book views in the Exchange 5.5 directory do not migrate because not all Exchange 2000 views can be represented in Exchange 5.5. Exchange 2000 allows better view control, based on LDAP filters.

Recipient Update Service

Each Exchange 2000 server updates the address lists by making calls to Wldap32.dll. Only one Recipient Update Service is active within each Active Directory domain; the others remain idle. The Recipient Update Service is fully integrated with the Exchange System Attendant (Mad.exe). According to the schedule you've set or by means of the **Update Now** option, the service contacts a local domain controller and proceeds to update address lists based on the rules set.

Offline Address Lists

To support mobile users, the offline address list allows a user to copy the contents of a server-based address book into a set of offline address book files (files with an .oab extension) on the local client's hard disk. The GAL is usually selected as the source for the offline address list generation.

You can configure offline address lists through System Manager. Under **Recipients,** right-click **Offline Address Lists,** select **New,** and then provide the offline address list name and list server. You can use any existing address list as an offline address list. There is a default offline address list created in the Exchange 2000 organization.

This offline address list can be associated with mailbox stores on the Exchange 2000 server, so that users with mailboxes on that server can download the files and work offline. To associate an offline address list with a mailbox store, right-click the mailbox store, choose **Properties,** and then, in **Offline address list,** click **Browse** and choose the offline address list.

When using Outlook 98 and Outlook 97 (version 8.03 and later), users are presented with a list of offline address lists to download. Outlook 2000 users can pre-select the offline address list. On the **Tools** menu, click **Settings,** and then select **Offline Folder Settings.**

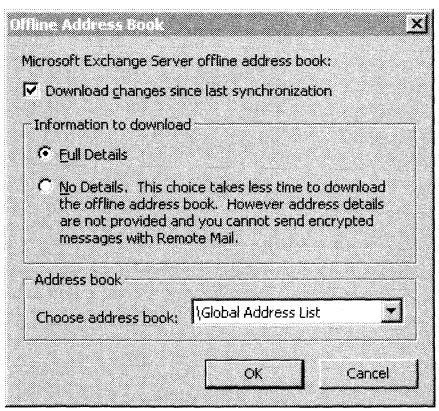


Figure 5.2 Offline address book settings

Client Access to Active Directory

More recent Outlook clients can access Active Directory directly. Other clients query Exchange when they need directory information. DSProxy provides referrals to the global catalog so that all clients can access Active Directory.

DSProxy Service

Outlook 98 (without the Quick Fix Engineering update) and earlier versions of Outlook use the DSProxy service to perform name resolution and address book searches against Active Directory. The DSProxy service mirrors Recipient Update Service in that it runs within the Exchange System Attendant (Mad.exe) rather than as its own service. The DSProxy service is an Exchange 2000 component that sends and refers MAPI directory service requests from Outlook clients to Active Directory for address book searches and name resolution.

DSProxy provides compatibility for clients with earlier versions of Exchange that are connected to an Exchange 2000 server. These clients assume that the storage and directory services are on the same server (as they were in Exchange 5.5), which may not be the case with Exchange 2000. If these services are not on the same server, clients that do not have Outlook 98 with the update patch or later versions of Outlook will not connect without the DSProxy service.

The DSProxy service connects to a global catalog server to send client requests back and forth. DSProxy enables Outlook clients to access the data within a global catalog, and isolates the clients from global catalog failure, because the DSProxy service always finds the nearest global catalog server by using DSAccess service (also part of System Attendant).

The DSProxy Name Service Provider Interface (NSPI) blindly forwards MAPI directory system requests to a global catalog server. Thus, the remote procedure call (RPC) packet is neither opened nor evaluated, because these actions would incur a significant Exchange 2000 server overhead, and compromise the security structure.

DSProxy begins by creating a listening thread on the domain controller for each supported network protocol, and it starts a single working thread for each processor. This can accommodate up to 25,000 client connections and dynamically adds more threads as required. A socket-mapping table keeps a reference of connections between clients and servers, ensuring that the correct responses from Active Directory pass to the associated client.

If the DSProxy server fails, DSProxy issues a callback to the System Attendant (Mad.exe), which in turn issues DSProxy a new target server name, known as re-targeting. The System Attendant will not re-target DSProxy without receiving a request.

DSProxy works over TCP/IP, Internetwork Packet Exchange (IPX), and AppleTalk protocols. However, it does not work over network basic input/output system (NetBIOS).

The DSProxy service performs the following functions:

- For earlier Outlook clients (Outlook 97 and Outlook 98 without the update patch), DSProxy sends directory requests on behalf of the clients to Active Directory by using the NSPI.
- For Outlook 98 with the update patch and Outlook 2000 clients, DSProxy provides a referral of the global catalog name directly to the client for smart MAPI clients.
- For all Outlook clients, DSProxy provides a layer of isolation and a referral service for Outlook clients (Outlook 97, Outlook 98, and Outlook 2000) in the event of a global catalog failure through querying the DSAccess service for the next closest global catalog server in the environment.

Upon startup, the Exchange System Attendant finds the closest domain controller in the domain by using the Domain Name System (DNS) resolver and passes the domain controller server name through to the DSAccess process (DSAccess.DLL).

The Active Directory domain controller used by the Exchange server is viewed in the properties of the Exchange server in the Exchange System Manager snap-in. See Figure 5.3. To view the domain controller, right-click the Exchange server, select **Properties,** select the **General** tab, and look in the **Domain controller used by services on this server** text box at the bottom of the tab.

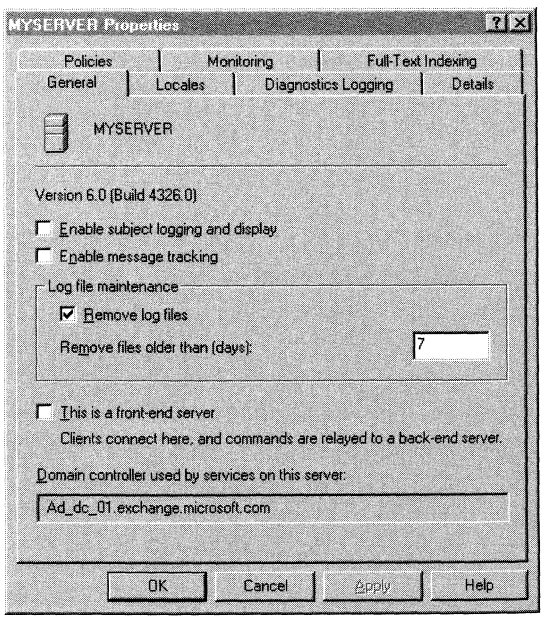


Figure 5.3 Exchange Sewer properties

Exchange Client, Outlook 7, and k 98

Earlier MAPI clients, such as the Exchange Client and Outlook 97 and Outlook 98, make MAPI directory service requests to an Exchange server. These requests range from resolving text in the **To** box of messages into user names, to showing objects from the global address list. To make Exchange 2000 compatible with the existing MAPI client base, an Exchange 2000 server proxy forwards any MAPI directory service requests through to a local global catalog server on the network. The DSProxy process on the Exchange 2000 server accomplishes this by forwarding a packet. It does not change the request into LDAP. Active Directory supports a number of protocols, including LDAP and MAPI directory service, so an Outlook directory request is valid, even when made directly to a global catalog server.

The global catalog server returns the result to the Exchange 2000 server, which in turn returns it to the requesting MAPI client. To the client, this process is seamless, and takes no extra time to complete.

The following steps are the communication process for a one-recipient name lookup:

Traffic and load generated by DSProxy:

1. The MAPI client sends one network packet to the Exchange 2000 server. The packet contains the search name in plain text.
2. The Exchange 2000 server sends the request to a local global catalog.
3. The local global catalog returns the result to the Exchange 2000 server.
4. The Exchange 2000 server returns the result to the MAPI client.
5. The MAPI client returns an acknowledgement to the Exchange 2000 server.
6. The Exchange 2000 server sends the acknowledgement to the local global catalog.

The directory search produces about six frames on the network. Even when multiple names are searched for in the directory, the name fragments go into a one-request packet.

If the user browses the GAL, the same process takes place. A few extra frames are produced on the network as the user scrolls through the address list, thus causing the client to retrieve more information.

Outlook 2000 and Outlook 98 SP2

Current versions of the Outlook client, including Outlook 2000 and Outlook 98 Service Pack 2 (SP2), use a slightly different method for address-book access. Initially, Outlook expects to find the directory service on the home Exchange server. Because the version of the Exchange server being used on the system can be determined only after loading Emsmdb32.dll (which is loaded after the address book provider, Emsabp32.dll), Outlook uses DSProxy for the first session. After the client contacts DSProxy (it will try all available transport protocols), a special referral passes back to the client, informing it that all future directory requests go to the global catalog server. Outlook saves this referral in the MAPI profile.

The referral mechanism seduces the load on the Exchange 2000 server and the latency for address-book searches. When an explicit server name goes into the profile, Outlook will need to be restarted if that Active Directory server fails. After restarting, the Exchange 2000 server passes a new referral to Outlook.

Coexistence and Upgrading

Microsoft has built several tools to help with the upgrade and coexistence of Exchange 5.5 with Exchange 2000 and Active Directory. In a coexistence scenario with Exchange 2000 Server and Windows 2000 Server, use the Active Directory Connector (ADC).

What is the ADC?

The ADC replicates and synchronizes Active Directory with an Exchange Server 5.5 directory. Because Active Directory is capable of holding all user account details, security and messaging information, a separate directory for Exchange 2000 is no longer necessary. The act of replicating Exchange 5.5 objects and attributes to Active Directory makes the objects and attributes available to Exchange 2000. Active Directory provides Exchange 2000 with global address lists, address book views, and offline address books.

Replication is the process of updating objects in a common directory namespace to one or more directories. Synchronization is copying objects from one directory to another with distinct namespaces. The ADC synchronizes changes in one or both directions, from Active Directory to the Exchange directory, from the Exchange directory to Active Directory, or both ways, depending on configuration. The ADC synchronizes two distinct types of directory objects:

- **Recipients** Mailboxes, distribution lists, custom recipients
- **Configuration information** Connectors, monitors, protocols, topology and other configuration information

The following are primary ADC features:

- Uses LDAP application programming interfaces (APIs) to replicate between two directories
- Hosts all active replication components on Active Directory, not on the foreign system
- Replicates only changes between the two directories, where possible
- Maintains object fidelity through replication (for example, matches an Active Directory group to an Exchange 5.5 distribution list)
- Hosts multiple connections on a single Active Directory server and manages them through connection agreements

When in native mode (a designated operating mode for an organization with only Exchange 2000 servers), Exchange 2000 Server enables separation of administrative and routing activities through administrative groups and routing groups, regardless of the underlying network structure. Note that during the migration, while coexisting with Exchange *5.5*, administration groups and routing groups are the same and appear as sites in the Exchange *5.5* directory. You can redefine routing groups once Exchange Server 5.5 is no longer in the environment.

ADC is installed in addition to Windows 2000 Server and Exchange 2000 Server. Upon installation, a Windows Service called the Active Directory Connector (ADC) appears. ADC can be started and stopped like any other service. An accompanying Active Directory Connector Manager console can be used to configure the relationships, called connection agreements, between Active Directory and the Exchange *5.5* target directory.

There are two versions of the Active Directory Connector:

ADC for Windows 2000

The ADC for Windows 2000 comes with Windows 2000 Server, and replicates directory information, such as mailboxes and distribution lists, from Exchange *5.5* directories to Active Directory and vice-versa. The Windows 2000 ADC is designed for customers who want to prepare Active Directory for Exchange 2000 during their Windows 2000 deployment and before installing Exchange 2000. For users of earlier versions of Exchange, much of the directory data can be uploaded to Active Directory at one time, thus decreasing installation time. Through synchronization, a specified Windows 2000 administrator can also manage Exchange *5.5* users.

Enhanced ADC for Exchange 2000

The enhanced ADC for Exchange 2000 is included with Exchange 2000 Server. The Windows 2000 ADC replicates *objects* in the Exchange *5.5* site (for example, the Recipients containers) to the Active Directory, whereas the Exchange 2000 ADC also replicates *configuration* data, such as protocol and connector data, and thus allows Exchange *5.5* and Exchange 2000 servers to coexist.

Because the Exchange 2000 ADC adds functionality beyond the Windows 2000 ADC, you should replace the Windows 2000 version with the Exchange 2000 version. The upgrade path between the two versions of the ADC is seamless. You may want to deploy the Windows 2000 ADC to populate your basic Active Directory quickly, and then install the Exchange 2000 ADC to migrate to Exchange 2000.

Site Replication Services and Recipient Update Service

Accompanying the ADC are several additional components installed by Exchange 2000 One of these is the Site Replication Services (SRS). The SRS helps maintain coexistence with Exchange *5.5* servers during the migration period. ADC, SRS, and Recipient Update Service are new services.

The Exchange 2000 SRS allows the Exchange 2000 system to emulate an Exchange *5.5* system (as perceived by the Exchange *5.5* part of the organization). SRS translates directory information between the Exchange *5.5* directory service and Active Directory. SRS incorporates the Exchange *5.5* intra- and inter-site replication engines to connect to Exchange *5.5* servers and the ADC. It has complete knowledge of the company's topology and adapts itself to any changes, preventing any interruption of service or loss of data.

The SRS is a modified Exchange *5.5* directory on the Exchange 2000 server that allows replication between Exchange 2000 SRS and remaining Exchange *5.5* servers as before. It is used as the Exchange 2000 side of the ADC configuration connection agreement. SRS is similar to Exchange *5.5* Directory Service, although MAPI is disabled.

The SRS enables you to avoid reconfiguring the endpoint of a connection agreement every time an Exchange *5.5* server in a site is upgraded to Exchange 2000 during the migration. It uses Exchange 2000 servers as bridgeheads for earlier Exchange sites through standard mail-based replication.

The Recipient Update Service, integrated with the system attendant, continuously polls the directory to determine which objects appear in the MAPI address book on global catalogs, and applies recipient policies (including proxy addresses). During polling, the service searches for all objects matching a set of specified rules. When Exchange 2000 is installed, an address list configuration object for that domain is generated; clients querying for address lists detect no difference. You can also use Recipient Update Service to set the *proxyAddresses* attribute for users, as well as other attributes, such as *homeMDB*, *homeMTA*, and *ms-Exch-Home-Sewer-Name*.

Why Use the ADC?

A temporary, but crucial step in a company's migration to Exchange 2000 is the Exchange 2000 ADC consolidation of an Exchange *5.5* mailbox with a Windows user account. This step creates a Windows 2000 mailbox-enabled user. The ADC is *required* in mixed-mode environments. Companies cannot use their own synchronization tools, because the ADC synchronizes information administrators may be unaware of, such as legacy distinguished names.

Active Directory must be fully populated with a user object for each Exchange *5.5* mailbox (with appropriate Exchange attributes filled out) before the first Exchange 2000 server, through upgrade or addition, is installed in the company. The Exchange 2000 server has no direct access to the Exchange *5.5* directory—the Windows 2000 Active Directory is the Exchange 2000 server's directory. Consequently, Active Directory must accurately reflect the information in the Exchange *5.5* directory, and the two directories must remain synchronized.

The ADC creates and maintains the bond between an Exchange *5.5* mailbox and a user until the last Exchange *5.5* server in the organization has been upgraded or decommissioned. Once the organization has only Exchange 2000 Server, the ADC is no longer needed.

There are four essential reasons for a company to install the Active Directory Connector:

- To leverage user information in the Exchange *5.5* directory and replicate the information to Active Directory to avoid re-entering the data.
- To enable Exchange 2000 installation while coexisting with an Exchange *5.5* environment.
- To create an environment where both Active Directory and Exchange *5.5* directory can be managed from one central place.
- To improve replication. In the Exchange *5.5* environment, Exchange server objects can only be modified in the Exchange site in which the object was created. In contrast, the Active Directory service allows any object modification on any server with a full replica of the Windows 2000 site, dramatically increasing replication flexibility. When directory objects change in Exchange Server *5.5,* the entire object replicates throughout the organization, because Exchange *5.5* supports only object–level replication. Active Directory employs a per-property replication (the changed property replicates, not the entire object), reducing replication time and replication conflict because modification of different attributes of the same object can occur in two locations simultaneously.

If you do not have at least one of the above requirements, reconsider installing the ADC. Exchange Server *5.5* can run without the ADC even when the domains and servers have been upgraded to Windows 2000 and Active Directory.

Connection Agreements

Installing the ADC on a server only defines a service within Windows 2000 and Active Directory. To establish a relationship between an existing Exchange site and Active Directory, configure a connection agreement. The connection agreement holds information such as the server names to contact for replication, object classes to replicate, target containers, and the replication schedule. It is possible to define multiple connection agreements for a single instance of ADC. Each one of these can go from Active Directory to a different Exchange site, or they can all go to the same Exchange site.

There are two types of connection agreements that you can view in ADC:

- User connection agreements replicate recipient objects and their data between the Exchange *5.5* directory and Active Directory. ADC displays these connection agreements. showing the direction and name of the connection agreement.
- Configuration connection agreements replicate configuration information, specific to Exchange, between the Exchange *5.5* directory and Active Directory. These agreements help Exchange 2000 to coexist with earlier versions of Exchange and Windows NT Server. ADC displays these connection agreements showing the bridgehead servers they use.

  Note The agreement for replicating configuration data is between Active Directory and SRS rather than Active Directory and an Exchange *5.5* server. In the following figure, the first connection agreement is the user connection agreement and the second is the configuration conneclion agreement.

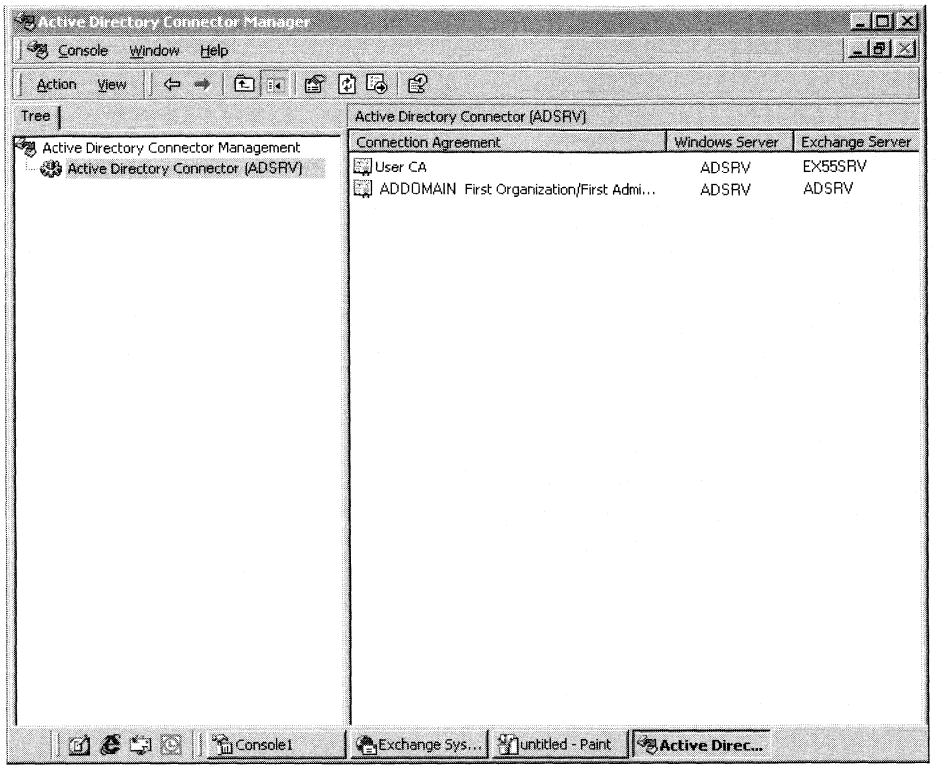


Figure 5.4 User connection agreement and configuration connection agreement

User Connection Agreements

Only user connection agreements that replicate directory information (recipient objects and their data) can be created with the ADC. Exchange 2000 creates configuration connection agreements automatically. User connection agreements are installed, configured, and controlled by the administrator.

A connection agreement defines the following:

- Directories to be synchronized.
- Windows 2000 Server synchronization objects.
- Exchange 5.5 synchronization objects.
- The direction in which synchronization takes place: one-way or two-way.
- Method for deleting objects.

- Synchronization schedule, which directly affects server-processing load. It is important that computers running the ADC have sufficient CPU and memory, and have high-bandwidth connections (LAN speed). The default synchronization schedule is eight hours, but can be adjusted for each connection agreement. For example, if you have multiple connection agreements to one Exchange *5.5* site, you can set one schedule to once per week, and another to once per day. For the initial ADC synchronization, you should set the connection agreements to different schedules to control the amount of data passing through and over the network.
- Advanced options, including:
  - Attribute mapping
  - New object creation
  - Authenticating directories
  - Specifying which organizational units or containers to synchronize

To perform two-way replication to centrally manage all objects, define a connection agreement to every Exchange site. This is to avoid directory triangulation problems. Because the ADC fully integrates Active Directory with the Exchange *5.5* directory, if it were possible to 'pull' Recipient containers from any Exchange site and 'push' these objects to any other Exchange site, with directory replication connectors deployed between the sites, it would cause a never-ending circular replication. To prevent this from happening, a connection agreement that is enabled for writing to the Exchange *5.5* directory can only interact with the objects in one Exchange site.

If you want to upload Exchange *5.5* directory data to Active Directory, configure a one-way connection agreement from Exchange. Because the ADC will not need to write information back to the Exchange *5.5* directory, recipient containers from multiple Exchange sites can be pulled through a single agreement. This is useful when you have many old Exchange sites in the company or when there is a distributed security model in place.

When defining a connection agreement to an Exchange site, the target directory bridgehead server must be running Exchange *5.5* Service Pack 3 or later, even if defining a one-way connection agreement. Other Exchange servers within the site are not required to run Exchange *5.5*.

Synchronization is simplified if you synchronize the entire Exchange *5.5* site instead of synchronizing individual recipient containers. Choosing the entire Exchange *5.5* site as the source and target of the connection agreement on the Exchange Server, and the Active Directory domain as the source and target on the Active Directory side of a two-way connection agreement effectively synchronizes the Exchange *5.5* recipient container hierarchy with the organizational unit hierarchy in Windows 2000. The organizational unit hierarchy or the location of individual recipients created by the ADC in the Active Directory can be changed later.

Configuration Connection Agreements

Kegardless of the connection agreements created for the Windows 2000 ADC (if present), the Exchange 2000 ADC creates a read-only configuration connection agreement that transfers bi-directional configuration information. Although the configuration connection agreement cannot be changed, the ADC writes changes to the appropriate directories.

When you install Exchange 2000 into an existing Exchange 5.5 organization, the first configuration connection agreement created by the SRS is called *Config CA_AdministrativeGroupName_Exchange2000ServerName*. The Exchange 2000 server must be in an Exchange 5.5 organization before you can view the configuration connection agreement. Although you can view the connection agreement in Active Directory Connector Manager, you can modify only a few settings. After replication, all Exchange 5.5 sites appear in Active Directory as administrative groups, and conversely, Exchange 2000 administration groups are represented in Exchange 5.5 as sites. This allows you to manage these objects from either side of the connection agreement. Although you can manage users in either Exchange 2000 or Exchange 5.5 using Active Directory Users and Computers, manage other objects (for example, mailboxes) in the administrator interface for the version of Exchange that contains the object.

Configuration connection agreements have a few attributes you can manipulate from the Active Directory Connector Manager, including:

- ***MsExchReplicateNow*** A Boolean value. If **True,** the ADC performs a replication cycle immediately for the connection agreement. Because this flag exists within the configuration naming context that is replicated around the forest, it is possible to set this value remotely. To set this from the Active Directory Connector Manager, right-click the configuration connection agreement and select **Replicate Now.**
- ***MsExchDoFullReplication*** A Boolean value. If **True,** the two directories are refreshed. To set this from the Active Directory Connector Manager, right-click the configuration connection agreement and select Properties, click the Schedule tab, and then select the **Replicate the entire directory the next time the agreement is run** check box.

In addition, configuration connection agreements have a number of attributes you can view by using a tool such as LDP.exe:

- ***msExchServer1HighestUSN*** The highest update sequence number found in the Active Directory. It is not recommended that you alter this attribute.
- ***msExchServer2HighestUSN*** The highest update sequence number found in the Exchange directory. It is not recommended that you alter this attribute.
- ***USNChanged*** The update sequence number assigned to the connection agreement object at last update. This attribute cannot be changed.

For more information about LDP.exe, see "ADC Object Matching" later in this chapter.

Creating Initial Connection Agreements

To retain the same container structure in Active Directory that you have in Exchange 5.5, create a single connection agreement for each domain, consolidating all Exchange containers into a single ADC–managed organizational unit in Active Directory. Once the objects replicate, use Active Directory Users and Computers to move those objects to the correct organizational unit. The ADC retains the relationship between the two objects through the object-GUID match, even though they have been moved.

Create one connection agreement for each Windows 2000 domain. Each connection agreement targets one Exchange 5.5 server in the Exchange site for that domain as a bridgehead server. Initially, you can deploy these connection agreements as one-way only, from the Exchange 5.5 directory to Active Directory, so that replication traffic is easily monitored, and object-attribute fidelity can be verified.

ADC Operation

When the connection agreement for the ADC is initially executed, the stored update sequence number of the target server is set to 0, and the ADC assutnes that this is the first time the agreement has been run. In addition, each connection agreement has its own signature, computed when you configure the connection agreement. Active Directory objects replicated into the Exchange 5.5 directory will have the signature of the earlier Exchange 5.5 bridgehead server, although the replication-signature attribute of the object will match the connection agreement's computed value.

When an Active Directory object replicates into the Exchange 5.5 directory, the *Object-Version* attribute will either be set to 1 if it is a new object or will be incremented by 1 if a modification is taking place. Before the LDAP write is made, the current value of *Object-Version* (if it exists) is read, incremented by 1, and written into the *Replicated-Object-Version* attribute on the Exchange directory object. If the ADC has just modified an object, both *Object-Version* and *Replicated-Object-Version* attribute are the same.

To detect changes in the Exchange 5.5 directory, an LDAP search of the Exchange 5.5 directory is performed with a request for objects that have an update sequence number between the value on the connection agreement (based on the last objects updated using the connection agreement) and the highest value in the Exchange directory. The search also returns tombstone objects.

> **Note** Tombstone entries are kept for 30 days (by default) before being deleted to leave enough time for them to be replicated throughout the company. Because deletions are changes on the local directory, they must replicate to all other Exchange 5.5 servers.

Additionally, a filter is set for the LDAP search to not request objects unless the *Object-Version* attribute is greater than the *Replicated-Object-Version.* This prevents the ADC from replicating objects originally sent to the Exchange directory back to Active Directory.

Active Directory uses attribute-based replication, unlike the Exchange directory, which uses object-based replication. To detect changes in Active Directory for replication to the Exchange environment, the connection agreement uses a combination of Active Directory update sequence numbers and the sum of attribute versions on the Active Directory object.

)C Object Matching

Exchange 5.5 recipient objects become native Windows 2000 objects in the Active Directory domain container. The following list describes how objects in Exchange 5.5 and Windows NT Server 4.0 correspond to Active Directory objects by default:

- **Exchange 5.5 Mailbox with a Windows NT 4.0 account** An Exchange **5.5** mailbox matches a mail-enabled disabled Windows user in Active Directory.
- **Exchange 5.5 Mailbox with a Windows 2000 account** An Exchange **5.5** mailbox matches a mailbox-enabled user in Active Directory; Mailbox-enabled users are security principals in Active Directory who can send and receive messages.
- **The *msExchHomeServerName*** This attribute contains the server name where the user sends and receives messages. Implicitly, this user type also has an SMTP address. The Active Directory object has its *legacyExchangeDN* attribute set to the distinguished name of the object in the Exchange Directory. All attributes from the two objects, such as telephone number, postal address, and so on, merge and populate to both directory objects.
- **Exchange 5.5 distribution list** An Active Directory mail-enabled group (type: distribution. scope: universal). Because the distribution list object may appear in Active Directory before the membership objects exist, these orphaned members will be binary encoded and written into the *unmergedAtts* attribute in the corresponding group entry. This ensures that membership changes to the Active Directory group object successfully replicate back to the Exchange directory if two-way connection agreements are in place. The objects listed in *unmergedAtts* are removed after these objects replicate. Windows 2000 uses groups to reduce the number of objects that require direct administration. Users are placed into groups, which can be further consolidated into other groups, allowing bulk administration. Exchange 2000 Server uses this concept of bulk administration and permission granting, but it also expands the use of groups by using them as distribution lists to which messages can be sent. Because Windows Groups are primarily used for security purposes (for example, to give a group permission to a file share or public folder), the Exchange distribution list becomes a new object class in Active Directory called Distribution-List. Active Directory supports two types of groups: security groups and distribution groups. Each of these can be scoped to domain local, global, or universal, and mail-enabled.
- **Exchange 5.5 custom recipient** A mail-enabled contact in Active Directory, that is, a contact with an e-mail address and (optionally) information that allows Exchange to properly route the mail. The Active Directory entry has only an SMTP address. An Exchange 5.5 custom recipient is typically a recipient using Lotus cc:Mail, Lotus Notes, or Microsoft Mail.

Table 5.6 shows Active Directory objects and the corresponding Exchange directory service objects.

Table 5.6 Matches between Active Directory and Exchange 5.5 objects

Active Directory Object	Exchange Directory Service Object (In the Target Container)
Mailbox-enabled user	Mailbox
Mail-enabled contact	Custom recipient
Mail-enabled user	Custom recipient
Mail-enabled group (type: distribution)	Distribution list
Mail-enabled group (type: security)	Distribution list

Note Active Directory objects that are not mail-enabled are not replicated back to the Exchange 5.5 server.

If at any point the ADC attempts to create an object already in the target directory, a numeric value prefixed with a hyphen appends to the common name of the object. The prefix ranges from –1 to –9999. This conflict resolution logic applies, in both directions, to objects in the Exchange directory or Active Directory. If the ADC finds that it's going to create a conflict with an existing object, it adds the -n to the end of the relative distinguished name of the new object to ensure it is unique.

ADC uses the set of target containers on the connection agreement to replicate directory objects that do not automatically map between the two directories. Configurable options for mapping objects include:

- Mailbox-enable an existing Windows 2000 user account.
- Mail-enable an existing Windows 2000 user account.
- Create a contact in the Windows 2000 container (mail-enabled or mail-disabled).
- Do not create any object.

By default, the ADC uses its own rules for matching attributes between the two directories. For example, the *Department* attribute in the Exchange directory is automatically matched to the *Department* attribute in Active Directory. Matching rules for all connection agreements reside in the default ADC policy object for the ADC in a field called msExchServerXSchemaMap, where X equals 1 for Active Directory-to-Exchange matches, and 2 for Exchange-to-Active Directory matches. You can view these rules by right clicking on **Active Directory Connector Management** in the Active Directory Connector Manager and choosing **Properties.** In some cases, it is desirable to change these rules. For example, you can match the Exchange *Custom-Attribute-3* attribute (called *Extension-Attribute-3* in LDAP) to the Active Directory *employeeID* attribute.

Matching can be done with a tool such as LDP, in binary mode, or through Active Directory Connector Manager. The LDP tool resides in the \Support\Tools subdirectory of the Windows 2000 Server installation CD. LDP, a graphical tool, enables LDAP operations, such as search, modify, add, and delete, in any LDAP-compatible directory. Administrators can use LDP to view objects stored in Active Directory, and object metadata.

Warning The rules format for matching attributes is very specific. You can use LDP to view attributes. However, without Administrator permissions, you may not be able to see all of the attributes of a particular object. For more information about LDP, see the Ldp.doc file in the \Support\Tools subdirectory of the Windows 2000 Server installation CD.

ADC Replication

Synchronous RPC is used for replication of the domain-naming context within a domain. When replicating data between two different systems, the data format and restrictions may be different. For example, the Active Directory supports UTF8 character format in the object name, but the corresponding distinguished name entry in the Exchange directory only supports teletext characters. Because the Active Directory schema contains formal definitions of every object class and its attributes, the ADC works out the restrictions imposed by the target directory and performs the necessary conversion. Schema conversion accommodates data discrepancies between the directories:

- Data format and presentation
- Field–length restrictions
- Mandatory and optional attribute mapping
- Single to multi-valued field mappings (and vice versa)

Replication network traffic from the Exchange 5.5 directory to Active Directory is approximately linear, about 140 KB to bind the directories, plus 14 KB per changed object. Table 5.7 gives some approximate results for the number of network frames and total network traffic sent from Exchange 5.5 to the ADC for different replications. (A frame is data that is transmitted between network points as a unit, complete with addressing and protocol information.)

Table 5.7 Network traffic for one-way replication

Number of Objects Replicated from Exchange 5.5 or SRS to ADC	Network Traffic
None	16 frames/3 KB
One	24 frames/8 KB
Two	27 frames/10 KB
Three	30 frames/12 KB

ADC Installation Location

For the best performance, install the Active Directory Connector on a member server in the Windows 2000 domain. Keep the replication schedule in mind, because the ADC can consume a lot of processor time when configured with multiple connection agreements. If you intend to install the ADC on a domain controller or global catalog, ensure that the server hardware can accommodate the extra processing load.

You should usually configure the connection agreement between a Windows 2000 global catalog server and the Exchange 5.5 server, or deploy the ADC in close network proximity to a global catalog in the same Windows 2000 site. In a multiple domain environment, you should have a global catalog server as the endpoint of every connection agreement, and have all global catalog servers for all domains within the same Windows 2000 site. If your Active Directory only consists of a single domain, there is no real difference between the objects and attributes held on the global catalog or a domain controller.

Multiple instances of ADC can be deployed in large installations, although you should configure the connection agreements so that different sets of objects are replicated by each connection agreement. Generally, multiple instances of ADC should not be deployed for fault tolerance, although you can reduce the risk of downtime by having more than one server running ADC. Note that although the connection agreements are defined and executed on the server running ADC, the source and target directories may reside on other computers.

Use directory replication bridgehead servers to facilitate Exchange 5.5 Server directory replication between Exchange sites. Where possible, use the same servers for ADC connection agreements. Ideally, these servers should be the last that you migrate to Exchange 2000.

In environments with multiple Active Directory domains, you should either point a separate connection agreement at an appropriate domain controller for each of the other Active Directory domains, or you should install ADC and the appropriate connection agreements on a separate server for each domain that you need to replicate with the Exchange 5.5 directory. Base your decision on load and network capabilities. For example, the ADC uses LDAP and RPC to communicate to the Exchange 5.5 directory, which requires direct Internet Protocol (IP) connectivity. Sometimes these protocols are blocked by a firewall, or the network transport may be incapable of supporting them. In such cases, install ADC on a server in each domain; the replication of Active Directory data between domains can take place over transports such as SMTP.

If possible, place the server with ADC on the same subnet as the Exchange 5.5 Server and Active Directory bridgehead servers.

Questions to Ask Before ADC Deployment

Do you want to manage all objects from the Active Directory?

The ADC enables you to use Active Directory Users and Computers to administer Exchange 5.5 mailboxes and Active Directory users. If this is a requirement, you need to deploy each connection agreement so that each agreement can write to the Exchange 5.5 directory. As the ADC replicates objects from the Exchange 5.5 directory to Active Directory, it must write various attributes into each Exchange object that it touches. Because of this, there is a one-time network performance hit when each Exchange object in the site is initially re-replicated within the native environment. In very large Exchange systems, this volume of changes can cause directory and network traffic problems. To reduce this risk, you can stagger the deployment of connection agreements.

If you can use the Exchange Administrator program to manage the environment of your earlier version of Exchange during the migration, most, if not all, connection agreements can be defined as One-way to Active Directory. In addition, you can deploy a single one-way agreement to one Exchange site, and acquire information from that point, instead of defining an agreement to each site.

How many Active Directory domains are you planning to have?

Unlike earlier versions of Exchange, Active Directory does not tie the namespace with the directory replication model. For example, when deploying Exchange 5.5, a company may need to seamlessly move users between Exchange servers. To meet this requirement, many create wide Exchange sites that span low-bandwidth networks. As the Exchange site forces the directory replication (and messaging) model to use RPC over slow network connections, additional management and administrative challenges occur.

Although Active Directory is far more flexible in terms of namespace and the replication model, Windows 2000 requires that domain-naming context replication (within the domain) occurs over synchronous RPC. In some environments, this may result in many small domains.

Another consideration for the ADC deployment is the number of servers running ADC that are required to replicate the data. Although an ADC can have connection agreements for multiple Active Directory domains, the protocol used between the ADC Server and Exchange 5.5 directory is mainly LDAP with a few RPC requests (RPC is for writing to the Exchange directory), and both require direct IP connectivity.

Will you upgrade your master accounts domain before deploying the ADC?

Each Exchange mailbox is matched to a primary Windows NT account. If the account's domain has been upgraded to Windows 2000 before the ADC is deployed, then you do not need to merge objects that exist in Active Directory with objects from the upgraded account domain.

If a domain has not been upgraded to Windows 2000, ADC creates new security principals to represent the users as mailboxes in the Windows NT 4.0 domain. ADC stamps the security descriptor so that the Windows NT 4.0 users can still access their mail. When the Windows NT 4.0 domain is upgraded, run Active Directory Account Cleanup Wizard to merge the ADC created objects with the upgraded Windows 2000 user accounts. Mailboxes are inaccessible to those accounts until this cleanup is done. The wizard is installed in the Exchsvr\Bin directory if you select the Microsoft Exchange System Management Tools option during Exchange installation, and is available from the Windows 2000 **Start** menu.

Assume that an existing Exchange *5.5* site has mailboxes that match primary Windows NT accounts in a domain containing only Windows NT Server 4.0. If a connection agreement is defined for the domain controller, ADC creates a mail-enabled disabled Windows user in Active Directory for the Exchange *5.5* mailbox. You can prevent the ADC from creating an object if a match cannot be established between the Exchange *5.5* directory object and the Windows 2000 Active Directory. When you upgrade the Windows NT 4.0 domain to Windows 2000, you must run the Active Directory Account Cleanup Wizard to merge these ADC objects with the new Windows 2000 user accounts.

How have you defined your container structure in the existing Exchange system?

By default, an Exchange *5.5* site has only the Recipients container defined. Some companies create other containers for different types of objects. It is uncommon for companies to have containers for each business unit, but special containers for holding different object classes, such as distribution lists and custom recipients, are not unusual. You may need to create multiple connection agreements to the same Exchange site.

For example, assume that an Exchange site has three containers: Recipients, distribution lists, and External. To replicate all objects from these containers to a single ADC-managed organization unit in Active Directory, create one connection agreement. However, if a similar organizational unit structure is proposed for Active Directory, define three connection agreements from the ADC Server to the Exchange site—one for each target container.

Preparing to Deploy the ADC

Exchange 5.5 Directory Preparation

For a smooth migration, clean up anomalous objects in the Exchange *5.5* directory before the first ADC replication. An anomalous object is defined as an Exchange *5.5* mailbox whose *mailNickname* attribute (Exchange *5.5 alias)* does not match the *loginID* attribute of the primary Windows security account associated with it in the same domain.

These mailbox objects or *resource* mailboxes are usually associated with a resource, such as a conference room, managed by several people over time. In other words, an Exchange *5.5* resource mailbox object with the *mailNickname* X may have its primary Windows user attribute set as Windows user *loginID* Y. However, the Windows user Y also has an Exchange *5.5* personal mailbox called *Y*, resulting in two Exchange *5.5* directory objects with the same primary Windows security principal.

In the Exchange 2000 and Windows 2000 environments, each mailbox can have only one primary Windows user account in the same domain associated with it. In the Exchange 5.5 environment this one-to-one association is not strictly enforced. In fact, the Exchange 5.5 mailbox is not required to have any association with a security principal, and it is possible for two or more to own the same Exchange 5.5 mailbox.

For example, Microsoft conference rooms have mailboxes, but the mailbox owner changes frequently. To illustrate, assume an e-mail name (and Exchange 5.5 mailbox name) is BldgZConf1 in the Exchange 5.5 global address list (e-mail address BldgZconfl @microsoft.com). BldgZconfl is associated with the Windows user account of an employee named Susan. Mail sent to reserve the conference room goes to Susan, who is currently associated with mailbox BldgZconfl . However, Susan also has her own mailbox.

If a mismatch is not cleaned up before the first ADC replication, the ADC can link the wrong primary Windows security principal to the mailbox. In the example, if the ADC were to encounter Susan's conference room mailbox first (Windows security principal), it would link the BldgZconfl mailbox to Susan as the primary owner. Meanwhile, as the ADC attempts to match Susan's Exchange 5.5 mailbox to the user object (Susan), it finds that it is already linked to the Exchange 5.5 mailbox object BldgZconfl. Susan then has no access to her personal mailbox, because she is no longer the primary owner (she is listed as primary for the BldgZconfl mailbox, and there must be a one-to-one correspondence). In addition, queries to the GAL for Susan's address and other information show the conference room's address, phone, and other information as Susan's personal information.

The same resource mailbox functionality can be achieved in the Windows 2000 and Exchange 2000 environment through the creation of disabled security principals (with no logon rights—for security purposes) in Active Directory, linked to each resource mailbox. When needed, include a number of different Windows security principals as secondary, or alternate, users of that account's mailbox, using the **Advanced View** tab of the Active Directory Users and Computers console. To create a simpler migration path, clean up the Exchange 5.5 directory in this way before first ADC replication.

Server and System Preparation

If possible, deploy two servers capable of supporting ADC. They should be brought online as member servers in the Windows 2000 domain. One acts as the main system, fulfilling the daily requirements of the ADC, while the other acts as a backup, providing redundant functionality in the event of a system failure. In the event of failure, the backup system goes online when the primary system becomes unavailable, maintaining ADC functionality until the primary system is repaired or rebuilt to act as the backup to the current system.

Accounts Preparation

The Windows account administrators can freeze changes to the Exchange 5.5 directory, during deployment, and alert the Accounts Group owners. This prevents ADC modification conflicts by the Exchange Administrator program and by Active Directory. However, a freeze may not be possible. Therefore, good communication between Windows account administrators and Exchange administrators is key to a successful installation.

Contingency Plan Preparation

In the event of a problem caused by ADC, perform the following incremental steps to get back to a known working state:

- Turn off the connection agreement.
- If this does not fix the problem, shut down the ADC service.
- If you have corrupt data, remove the organizational unit where the data was stored, and perform a restore from a tape with a known good state.
- As a last resort, use the Windows 2000 LDP tool to edit data in Active Directory. For more information about using LDP, see "ADC Object Matching" earlier in this chapter.

Exchange Organization Preparation

Perform the following steps before running Exchange 2000 Setup to prepare Active Directory for mixed-mode environment during migration:

- Extend the Active Directory schema with schema and display specifiers.
- Create groups.
- Delegate authority.
- Delegate predefined domain level permissions to specific Exchange 2000 processes, such as Recipient Update Service.
- Perform explicit access control list (ACL) modifications to prevent current domain and enterprise administrators from obtaining Exchange 2000 administrator permissions through inheritance.

This last step ensures that the new administrative structure can be created during migration.

To accomplish these steps in an environment with Active Directory, Exchange 5.5, and Windows domain administrators have differing permissions, you must use two setup utilities, ForestPrep and DomainPrep. ForestPrep is run once per forest to extend the Active Directory schema, nominate the Exchange 2000 administrators, create the Exchange 2000 organization object in the configuration naming context, and set up the permissions structures. DomainPrep is run once per domain to create the public folder proxy container and set permissions within the domain. You do not need to run DomainPrep in a domain until you are ready to install Exchange 2000. DomainPrep must be run in all domains where you install Exchange 2000 and in all domains that contain recipient objects, such as mailboxes or distribution lists.

ForestPrep and DomainPrep prepare an organization for Exchange 2000 Server installation by distinguishing the parts of setup that require high-level network permissions from those that only need Exchange administrator permissions. You do not need to run ForestPrep and DomainPrep if all the following conditions are met:

- All Exchange 2000 servers will be in a single domain.
- The domain contains the schema master.
- All Exchange users are in the domain.
- The account installing Exchange 2000 has Enterprise Admins and Schema Admins permissions.

To run ForestPrep, gather the following information and perform the following tasks:

- Know what account will be the Exchange 2000 Administrator account.
- Know whether this is going to be a new Exchange 2000 organization or a connection to an existing Exchange 5.5 organization.
- Know what you will name the organization if it is new.
- Know the name of the existing Exchange 5.5 server and the Exchange 5.5 service account password, if you are joining an existing Exchange 5.5 organization.
- Run ForestPrep in the same domain as the Schema Master.
- Run ForestPrep from an account with Windows 2000 Enterprise Admins and Schema Admins permissions.
- If you are joining an existing Exchange 5.5 organization, run ForestPrep from an account that has at least View-Only Administrative permissions on the Exchange 5.5 site and configuration container.

If you are joining an existing Exchange 5.5 organization, install Exchange Server 5.5 Service Pack 3 and the Active Directory Connector that shipped with Exchange 2000 before you run ForestPrep. The Active Directory Connector that shipped with Windows 2000 is insufficient.

ForestPrep nominates the first full Exchange 2000 administrator account. This account assigns delegated Exchange administrative rights to other new Exchange administrator accounts using the Exchange Administration Delegation Wizard. When ForestPrep and DomainPrep have finished, this account will be equivalent to an organizational-level Exchange full administrator. Once this account has been added to the local Administrators group, it can install the first instance of Exchange 2000 (without belonging to Enterprise, Schema, or Domain Admins security groups). This account can also delegate various Exchange 2000 administration roles to other users or groups, by using the Exchange Administration Delegation Wizard. ForestPrep will:

- Extend the Active Directory schema to include Exchange-specific information. This affects the entire forest and typically takes about a half hour. If the forest is large, it may take considerably longer to replicate changes to every domain and domain controller.
- Prompt for and create the Exchange organization name and object in Active Directory. After you run ForestPrep and then run DomainPrep in each domain where Exchange is to be installed, Exchange 2000 Setup queries Active Directory for configuration information. This simplifies the deployment of Exchange throughout the forest.
- Assign Exchange Full Administrator permissions for the account that you specify. This account has the authority to install Exchange throughout the forest. After the first installation of Exchange 2000, use this account to run the Exchange Administration Delegation Wizard, which configures Exchange-specific roles for administrators throughout the forest.

To run ForestPrep

1. Insert the Exchange 2000 Server compact disc into your CD-ROM drive.
2. On the **Start** menu, click **Run** and type ***E*:\setup\i386\setup /ForestPrep**, if **E** is your CD-ROM drive. Spell ForestPrep correctly (it is not case-sensitive). Otherwise, setup runs and ForestPrep will not.

DomainPrep is run in every domain with recipients and in every domain where Exchange 2000 will go. In the latter case, run DomainPrep after ForestPrep, and after replication occurs, in every domain where an Exchange 2000 server is going to be installed. It takes about 15 to 20 minutes for a parent and child forest in a single Windows 2000 site, but can take longer with grandchild domains or domains in different Windows 2000 sites. You can check replication progress with the Replication Monitor tool in Windows 2000.

Run DomainPrep in each domain where you want to install Exchange 2000, including the same domain where you ran ForestPrep. DomainPrep can be run from any server in that domain. You must also run it in domains where Exchange 2000 is not installed, but where you want Exchange 2000 mail-enabled users for whom you want to create a Recipient Update Service.

DomainPrep creates the necessary global groups for Exchange administration and directory service to Active Directory. The account for this utility must have Domain Admins and local Administrator permissions in each of the domains to which the ADC will connect. Exchange 2000 and Exchange 5.5 Administrator permissions are not required to run DomainPrep.

DomainPrep will:

- Create two new domain groups: the global security group, Exchange Domain Servers (formerly known as Domain EXServers), and the domain local security group, Exchange Enterprise Servers (formerly known as All Exchange Servers).
  - Exchange Domain Servers contains the computer accounts of all the Exchange 2000 servers in the domain and is needed for Recipient Update Service.
  - Exchange Enterprise Servers contains the Exchange Domain Servers groups from all the domains running Exchange (or where DomainPrep has run or where there is an active Recipient Update Service).
- Create the public folder proxy container.
- Grant appropriate permissions for Exchange 2000 Administrators and Exchange Servers on these objects.

To run DomainPrep

1. Insert the Exchange 2000 compact disc into your CD-ROM drive.
2. On the **Start** menu, click **Run** and type ***E*:\setup\i386\setup /DomainPrep**, where **E** is your CD-ROM drive. Spell DomainPrep correctly (it is not case-sensitive). Otherwise, Setup runs but DomainPrep does not.

Minimum Requirements for Installing the ADC

Install the ADC on a Windows 2000 member server, meeting or exceeding the minimum requirements for an Exchange 2000 server. They are:

- Intel Pentium 133–MHz or faster processor
- 128 MB of RAM (256 MB recommended)
- Microsoft Windows 2000 Server operating system
- At least 500 MB of available disk space for Exchange 2000 on the installation drive
- At least 200 MB of available disk space on the system drive
- Page file set to a minimum of twice the amount of system RAM (recommended)
- VGA-resolution monitor, or higher
- A CD-ROM drive

In addition, the following requirements must be met:

- The Windows account used for ADC installation is not an Exchange 5.5 service account.
- The Windows account used for installation is logged in to the same domain as the server.
- The Windows account used for installation has access to the global catalog (must be a member of the Schema Admins Security Group).
- The Windows account used for installation has Domain Administrator rights for each domain the ADC will connect to (member of the Domain Admins security group) and for the member server on which the ADC is installed.
- The ADC service account specified during Setup is a member of the Enterprise Admins or root Domain Adrnins group. This successfully updates user information in all necessary domains.
- Windows 2000 and DNS services are deployed in accordance with Windows 2000 documentation. Accurate execution of DNS or dynamic DNS services in the Windows 2000 environment is critical to ADC (and Exchange 2000) functionality.
- The ADC-target Exchange 5.5 server has Service Pack 3 installed.

  Note Separate servers into dedicated service roles for manageability and reliability. Services, such as SMTP, located on the same server with ADC will present additional configuration complexity.

Monitoring the ADC

The ADC records several categories of events in the event log's application section. You can change logging levels and select logging categories for ADC. Each logging category provides informational, warning, and error messages.

To configure diagnostic logging

1. In the Active Directory Connector Manager MMC, right-click **Active Directory Connector,** and then click **Properties.**
2. On the **Diagnostics Logging** tab, under **Category,** select the category to log, and then select a **Category Logging Level.**

The following table shows some of the ADC event categories.

Table 5.8 ADC events that can be logged

Category	Instances of Events Messages
Replication	Shows events that occurred during replication.
Account management	Shows events that occurred while attempting to write or delete an object.
Attribute mapping	Shows messages indicating events that occurred while matching attributes between Active Directory and the Exchange directory.
Service Controller	Shows events that occurred while the service is started or stopped.
LDAP operations	Shows events that occurred while accessing the directory using LDAP.

Because the log files are not deleted by default, you can leave the logging level set to **None** (the default) unless you are actively troubleshooting an issue and need to monitor events. The following table summarizes the most important application event identifiers (IDs) monitored.

Table 5.9 Important ADC event IDs

EventID	Description	Explanation	Action
8084	Could not load mapping table.	Active Directory Connector could not load mapping table.	Verify that the ADC account has permission to read the connection agreements and all of their attributes.
8113	The service could not be initialized because the necessary file X could not be found. Make sure that the operating system was installed properly.	Active Directory Connector service initialization failure: missing file.	Check the file version and verify that its location is in an executable path.

Table 5.9 Important ADC event IDs *(continued)*

8114	The service could not be initialized because the necessary entry point X could not be found in the file Y. Make sure that the operating system was installed properly.	Active Directory Connector service initialization failure: entry point could not be found.	Check the file version and verify that its location is in an executable path.
8141	The operating system has run out of resources. The connection agreement will shut down and then restart. If this problem persists, try restarting the ADC or the server.	Active Directory Connector connection agreement shutdown and restarted.	Other applications running on the system share resources. This is a Windows-level error that indicates a shortage of resources for all applications. To reduce the load on resources, restrict unnecessary applications from running on the server.
8142	Active Directory Connector unexpected exception.		Restart the service. Verify that there is enough disk space and memory.
8143	Active Directory Connector out of memory exception.		Increase the swap file size, close some other applications, or restart the service.
8144	Exception X was raised at address Y.	Active Directory Connector exception.	Restart the service. Verify that there is enough disk space and memory.
8204	Error loading import rules	Active Directory Connector error loading import rules.	
8223	ADC cannot replicate X because too many of its attributes conflict with other objects.	Active Directory Connector unable to replicate.	
8002	The service was stopped.	Active Directory Connector service stopped.	The Service Control Manager has returned this request. The service was shutdown successfully. This message requires no action and is informational only.
8026	Connection is down.	The domain controller with the connection agreement is down.	

Deployment Tips

- Allow sufficient time after you run ForestPrep, DomainPrep, and ADC installation for complete replication of Active Directory changes throughout the environment. When ADC first runs, the ADC replicates every object not in Active Directory. All those objects then replicate to every domain controller and global catalog in the Exchange organization. This may take some time, especially in larger environments. After the first peak of replication traffic, only ongoing incremental changes are replicated, until the last Exchange 5.5 server is removed or upgraded.
- Because the ADC never really finishes during the migration (it continuously keeps the two directories synchronized), consider installing the ADC, SRS, and Recipient Update Service on a dedicated Windows 2000 member server for ease of management and increased availability.
- Understand the security implications of groups. Because Active Directory groups are used for both access control and distribution lists, be sure to prevent modifications to group membership (when the ADC maps Exchange 5.5 distribution lists to Active Directory groups) that compromises system security.
- Understand the differences between connection agreement directions. For example, the results of two one–way connection agreements may differ from the results of one two–way connection agreement. Carefully configure the **From Windows** and **From Exchange** tabs of Active Directory Connector Manager, and the source and target containers. For the desired results, match both the correct containers and the direction of information flow.

CHAPTER 6

Deployment Strategies

Peter Nilsson, Principal Consultant, Microsoft
Will Martin, Senior Consultant, Microsoft

A good deployment strategy saves time, money and frustration. Use this chapter to examine the issues involved in developing a deployment strategy before you begin.

Every company requires a different approach to upgrades, migration, and deployment. The discussion in this chapter can help you decide which approach you should take to achieve your goals with the least disruption, at the lowest cost. Although this chapter doesn't include details about costs, effective planning can prevent the need to revisit and re-engineer at a later date—which keeps your costs down.

To help you build a highly effective Microsoft Exchange 2000 infrastructure for routing and managing e-mail messages, this chapter focuses on high-level planning issues. It does not provide detailed procedures describing how to perform tasks such as upgrading different types of connectors.

In This Chapter

Deployment Roadmap

Preparing to Deploy

Populating User Accounts in Active Directory

Upgrading Windows NT Server to Windows 2000 Server

Upgrading Exchange Mailboxes

Scenarios

Deployment Roadmap

Most of this chapter discusses the deployment of Exchange 2000 in an existing Exchange 5.5 organization. Many of the considerations for this situation do not apply if you are deploying Exchange for the first time. You can concentrate solely on issues such as whether your design should be centralized or distributed, how many users to support per server, and how large you can allow databases to grow.

New Exchange 2000 Organization

If you are creating a new Exchange 2000 organization, the information in the rest of this chapter can still be useful, especially if your design is complicated by either of the following factors:

- The existence (planned or real) of multiple forests in your organization.
- The need to merge an Exchange 5.5 organization into the Exchange 2000 network at some point in the future.

Because Exchange 2000 is not restricted by Exchange 5.5 site topology, you can deploy servers in any order, and in almost any design. You can easily modify the topology to fit changing circumstances.

The key planning issue that is new for Exchange 2000 regards effective integration with Active Directory; for example, ensuring correct placement and availability of global catalog servers.

Existing Exchange Organization

Upgrading to Exchange 2000 usually involves four distinct phases, as outlined in Table 6.1. Each phase requires planning and is described in detail later in this chapter.

> **Note** Three concepts in Table 6.1 are discussed in "Upgrading Exchange Mailboxes," later in this chapter. The in-place upgrade method involves performing an upgrade on a single functioning Exchange 5.5 server. The move mailbox upgrade method involves moving mailboxes from an Exchange 5.5 server to an Exchange 2000 server. The leapfrog upgrade method involves using the move tnailbox upgrade method in a specific sequence.

Table 6.1 Deployment roadmap

Phase	Available Methods
Preparation	Run ForestPrep and DomainPrep.
Populate Active Directory	Perform in-place domain upgrade, followed by running ADC (Active Directory Connector).
	Run ADC, upgrade later and run Active Directory Cleanup Wizard.
	Run ADC, clone user accounts later by using Active Directory Migration Tool.
	Run Active Directory Migration Tool with SIDHistory, and then run ADC.
	Run Active Directory Migration Tool without SIDHistory, recreate Access Control Lists (ACLs), and then run ADC.

Table 6.1 Deployment roadmap *(continued)*

Phase	Available Methods
Upgrade the operating system for Exchange servers to Microsoft Windows 2000 Server	Upgrade existing resource domain (in which Exchange is running) to Windows 2000 and connect to existing forest.
	Do not upgrade if computers running Exchange 5.5 will be retired or recycled.
	Upgrade existing servers and move them to a new Windows 2000 domain.
Upgrade or migrate existing Exchange mailboxes	In-place upgrade.
	Move mailboxes to new or upgraded servers.
	Perform leapfrog upgrade method.

With this roadmap, each deployment phase can be planned largely in isolation from the others. Note that the different phases are dependent on each other, and that the overall system architect should have a good grasp of the interdependencies. Separate teams can be assembled to plan the different stages. These deployment teams can make the best use of specialized skills that exist and compensate for skill gaps within the organization.

After initial deployment decisions have been made, it should be necessary to concentrate on only one stage at a time. This streamlines the overall planning and design process.

Preparing to Deploy

The first phase of the Exchange upgrade is small but important for the Active Directory administrator.

Although it isn't necessary to perform this phase separately for technical reasons, you may want to separate it based on your administrator constraints.

Two processes need to be complete before Exchange 2000 can be installed on a server in the forest:

- The Active Directory schema must be updated.
- Appropriate permissions must be assigned.

Exchange 2000 Setup is designed to allow these two processes to be executed separately from the installation or upgrade of a server. Run ForestPrep to apply schema changes and run DomainPrep to set appropriate permissions. For more information, see *Microsoft Exchange* 2000 *Sewer Planning & Installation.*

Deciding Whether to Install ADC First or Run ForestPrep

As you prepare for a system deployment, you can install ADC either before or after you run ForestPrep. Depending on whether you are installing a new Exchange organization or joining an existing Exchange 5.5 organization, ForestPrep presents different options.

If you are installing Exchange 2000 for the first time, complete the following steps in the order given:

To Install Exchange for the first time

1. Run ForestPrep.
2. Run DomainPrep.
3. Install Exchange 2000.

If you are upgrading an existing Exchange 5.5 organization, you must specify this when you run ForestPrep. Though it is not necessary to configure connection agreements before you run ForestPrep, ADC must be installed.

To upgrade an existing Exchange 5.5 organization

1. Install ADC.
2. Run ForestPrep. Choose to join an existing Exchange 5.5 organization. ForestPrep then contacts the Exchange 5.5 directory that you specify and builds a root Exchange 2000 organization in Active Directory.
3. Run DomainPrep.
4. Run Exchange 2000 Setup either to upgrade an existing Exchange 5.5 server, or to add a new Exchange 2000 server to an existing Exchange 5.5 site.

Completing the preceding steps causes the following to occur:

- After you install ADC, a number of schema changes are made and replicated to all domain controllers in the forest. In addition, because the partial attribute set is changed, all global catalog servers replicate partial replicas from other domains. Several additions to the configuration naming context also replicate to all domain controllers in the forest.
- After you run ForestPrep, more schema changes are made and replicated to all domain controllers in the forest. The partial attribute set does not change, so global catalogs do not replicate again. Some additions are made to the configuration naming context and these additions are also replicated to all domain controllers.
- After you run DomainPrep. small changes are made, primarily to the local domain naming context. These changes are replicated to all domain controllers in the local domain.
- After you run Setup, when the first Exchange 2000 server is installed, a configuration connection agreement is created and information about the Exchange 5.5 network is written to the configuration naming context. This is replicated to all domain controllers in the forest.

In addition, before anyone can use Exchange 2000, you must run user connection agreements to synchronize Exchange 5.5 mailboxes and Windows 2000 user accounts. This may cause additional information to be replicated within affected domain naming contexts, and also to global catalog servers. Only attributes that change replicate to global catalog servers; the entire objects do not replicate.

An Exchange 2000 server must be placed close to a global catalog server on the network, usually on the same LAN segment (or equivalent). In addition, global catalog servers must have the hard drive space and RAM to handle the demand that will be generated by Exchange.

It is important that you understand and plan for the frequency and volume of directory changes that occur after Exchange 2000 becomes an integrated part of your infrastructure. If you have deployed any earlier versions of Exchange, the administrators of the earlier versions can help you to predict the number and types of directory changes. Your Active Directory infrastructure must be designed to handle the load of both Windows 2000 Server and Exchange 2000 Server.

Populating User Accounts in Active Directory

This chapter assumes that you are upgrading an existing Exchange 5.5 installation. In an organization with more than one Exchange server, this means that a coexistence phase is necessary when Exchange 5.5 and Exchange 2000 servers run as part of the same Exchange organization. Because Exchange 2000 no longer has its own directory and relies entirely on Active Directory, all of the Exchange 5.5 users need to be represented as mailbox-enabled users in Active Directory before the first Exchange 2000 server can be a working part of the Exchange organization. Users must be represented as mailbox-enabled users complete with all the configuration details that Exchange 2000 needs to address and route mail correctly.

The only way to get the correct Exchange data into Active Directory is to use ADC. The role of ADC is complicated by the fact that you are not just upgrading Exchange, but synchronizing that upgrade with an upgrade or migration from Microsoft Windows NT version 4.0 to Windows 2000 Server. Upgrading to Exchange 2000 and synchronizing with migrations or upgrades to Windows 2000 Server can be complex and can occur in various phases. However, before installing ADC, you should complete the main tasks to upgrade or migrate to Windows 2000 Server.

Review this section to see what upgrade or installation tasks can be done at the same time as installing ADC. All other upgrade and migration tasks should be completed before installing ADC. This section describes ways to synchronize the actions of ADC with what occurs during the Windows NT 4.0 domain upgrade.

In practice, installing, configuring, and running ADC can be the first step of upgrading your first Exchange server. But because this step is so complex, the planning is discussed independently in the following scenarios.

Scenarios

This list describes five scenarios for using ADC to migrate accounts to Windows 2000:

- Perform an in-place domain upgrade, and then run ADC.
- Run ADC, and then upgrade later and run Active Directory Cleanup Wizard.
- Run ADC, and then clone later using Active Directory Migration Tool.
- Run Active Directory Migration Tool with SIDHistory, and then run ADC.

  Run the Active Directory Migration Tool without SIDHistory, recreate Access Control Lists (ACLs), and then run ADC.

This section describes each scenario and the best way to achieve the same goal in each case: Create a single Active Directory user account with all the correct Exchange attributes for each Exchange user who has a Windows NT 4.0 account.

In-Place Domain Upgrade Followed by ADC

An in-place domain upgrade is a straightforward scenario. To begin with, all the Windows NT 4.0 account domains (all the domains that contain accounts that have been used to access Exchange mailboxes) must be upgraded to Windows 2000. At a minimum, the primary domain controller in each relevant domain must be upgraded. All the domains must be in the same forest: A maximum of one of the upgrades could have created a new forest and the rest must have joined their domains to the existing forest either as child domains of an existing domain or root domains of a new tree.

You can now install and configure ADC with connection agreements that synchronize Exchange 5.5 mailboxes with Active Directory user accounts. ADC matches the security identifier (SID) from the primary Windows NT Account with the primary SID of an Active Directory user account. A SID is statistically unique number that identifies each user and group. When a new user or group is created, Windows 2000 generates a SID for the account. The operating system uses the identifier to verify access permissions when a user requests access to an object, instead of using the user name. Because the SIDs did not change during the upgrade from Windows NT 4.0 to Windows 2000, ADC matches the mailbox to a user account based on the existing SIDs. Then, ADC populates the existing account in Active Directory with information from the Exchange 5.5 directory, such as an e-mail address, a phone number, or an office location.

Run ADC, Upgrade Later, and Run Active Directory Cleanup Wizard

You may want to start upgrading to Exchange 2000 before you migrate to Windows 2000 Server. If so, you can use ADC to populate Active Directory with the basic user accounts that are required for Exchange 2000. Then at a later date, you can upgrade existing Windows NT 4.0 account domains, and then merge the resulting duplicate account objects in a clean-up stage.

Before running ADC, a basic Active Directory infrastructure must be in place. In addition, you must specify a target domain in which ADC creates disabled user accounts. For the purposes of this discussion, assume that the target domain is the root domain of the forest.

When ADC runs, a disabled user account is created in Active Directory for each Exchange 5.5 mailbox. Each disabled user account contains all of the Exchange configuration information, such as e-mail addresses and the name of the private information store that contains the mailbox. Mailbox permissions are copied from the Security Descriptor of the Exchange 5.5 mailbox and assigned to the Windows 2000 user account.

When ADC is run in this way, the *msExchMasterAccountSid* attribute is created on each disabled user account and assigned the SID from the primary Windows NT 4.0 account for the Exchange 5.5 mailbox.

Everything is now in place to begin the upgrade to Exchange 2000. You can upgrade user mailboxes to Exchange 2000 either by moving the mailboxes to an Exchange 2000 server or by upgrading the server on which the mailboxes reside. Users can continue to access their mail by using their Windows NT 4.0 security credentials.

You will upgrade the Windows NT 4.0 account domains to Windows 2000 later. At this point, these domains must join the forest that Exchange 2000 is using, either as child domains of the existing root domain or as root domains of a new tree. There are now two accounts in Active Directory for each Exchange user: the original disabled user account and the newly upgraded account.

> **Important** Exchange 2000 cannot be used while these duplicate accounts exist in Active Directory.

The solution is to run Active Directory Account Cleanup Wizard, which is included with Exchange 2000. Active Directory Account Cleanup Wizard goes through Active Directory to match and merge duplicate accounts. The wizard does this by matching the *msExchMasterAccountSid* attribute created on each disabled user account by ADC with the primary SID of the upgraded user account.

Active Directory Account Cleanup Wizard merges all the attributes of the disabled user account into the upgraded user account. As part of the process, the disabled user account is deleted, leaving a single mailbox-enabled user account in Active Directory for each Exchange user.

Run ADC and Migrate Later By Using the Active Directory Migration Tool

This scenario starts off in the sane way as the previous one, except that there is no intention of upgrading the Windows NT 4.0 account domains; instead, you will migrate user accounts later.

ADC creates the disabled user account in the same way as described in the previous scenario. You can then begin the Exchange 2000 deployment.

To migrate users, select one of the following methods:

- Merge accounts during migration.
- Migrate to a different account and merge later (by using Active Directory Cleanup Wizard).

The first method is preferred; with the second method you create duplicate objects, replicate them around the entire Active Directory. then merge details in a second operation that generates another round of replication.

You can use tools like Active Directory Migration Tool to merge accounts during migration. If you use Active Directory Migration Tool to migrate accounts into the domain that contains the disabled user accounts created by ADC, Active Directory Migration Tool detects the duplicate accounts and follows whatever conflict handling behavior has been chosen. The options for conflict handling are *ignore,* which does nothing; *replace,* which merges accounts; and *rename,* which creates a separate account using a prefix or suffix supplied by the administrator. When you merge accounts during migration, you must choose the replace option when you run Active Directory Migration Tool.

Note Third-party tools generally offer equivalent functionality.

The second method—to migrate to a different account and merge later—is more complex; how you proceed depends in part on how you migrate the accounts.

If the migrated accounts are cloned (that is, the sIDHistory attribute is populated with Windows NT 4.0 SIDs), the matching required for the merge is easy because the Active Directory Account Cleanup Wizard looks in sIDHistory if it doesn't find a match on the primary SID. You can clone accounts only if the target Windows 2000 domain is in native mode.

If accounts are not cloned (that is, they are migrated without putting anything into sIDHistory), the Active Directory Account Cleanup Wizard cannot match old accounts with new accounts on its own. In this case, you can still manually associate accounts; the Active Directory Account Cleanup Wizard finishes the merge operation. However, merging a large number of accounts by using this method would be tedious and time-consuming, with the possibility of introducing errors.

The second method may also affect the way you use domains during the upgrade or migration process. In Exchange 5.5, when users and mailboxes are created together using the Exchange Administrator program, the Windows NT 4.0 user accounts and Exchange directory names (RDNs) are the same by default. If they are the same, you cannot migrate those user accounts directly to the domain that contains the disabled user accounts that were created by ADC. You must specify that the disabled user accounts not be created in the domain to which you want to migrate the users.

Therefore, it is recommended that you create a child domain under your root domain and use this as a *transition* domain in which to create the disabled user accounts. This allows you to migrate user accounts directly to the domain you want. You can remove the transition domain after the Active Directory Account Cleanup Wizard has deleted all the disabled user accounts.

Run Active Directory Migration Tool with sIDHistory and Then Run ADC

Running ADC after fully migrating all your user accounts is simple. Although you can use Active Directory Migration Tool to clone users from one or more Windows NT 4.0 source domains into one or more Windows 2000 domains, you should try to consolidate user accounts. In this scenario, you configure Active Directory Migration Tool so that the sIDHistory attribute is updated with primary and group SIDs from the source account. Cloned users get new SIDs, but they can still access resources for which permissions have been assigned by using their old SIDs.

The key is that ADC recognizes sIDHistory and can match the Exchange 5.5 mailbox's primary Windows NT 4.0 account with the old primary SID that is now stored in sIDHistory.

Run Active Directory Migration Tool Without sIDHistory and Then Run ADC

Not everyone who migrates to a new Windows 2000 environment needs to clone users. Security and administration issues with sIDHistory may cause you to decide to migrate without copying users' old SIDs.

A newly migrated account without sIDHistory cannot access any of the resources to which the old account had access. For this reason, the migration process must include steps to reapply the new SID to all of the resources to which the old SID had been granted permissions. To avoid a situation in which users cannot access resources, perform this step before changing the user over to the new account.

For this approach to work well with ADC, you must update permissions on Exchange 5.5 mailboxes before you run ADC. It is easy for ADC to match mailboxes and accounts after you have updated the primary Windows NT 4.0 account on each Exchange mailbox to reflect the new SID of the user's Windows 2000 account.

You can update permissions manually; however, this is not realistic for large numbers of users. Active Directory Migration Tool includes an option to recreate Access Control Lists (ACLs) that set permissions on Exchange mailboxes. You can use Active Directory Migration Tool or third-party domain migration tools to recreate ACLs that are associated with migrated accounts.

Upgrading Windows NT Server to Windows 2000 Server

Exchange 2000 runs only on a computer running Windows 2000 Server that is a member of a Windows 2000 domain (mixed mode or native mode). Because you are starting with a network of Windows NT 4.0 servers running Exchange 5.5, you must upgrade Exchange servers from Windows NT 4.0 to Windows 2000 Server before you can upgrade from Exchange Server 5.5 to Exchange 2000 Server.

For this discussion, refer to the third phase of the roadmap: upgrading servers to Windows 2000 Server. You have three options:

- Do not upgrade your entire network.
- Upgrade an existing resource domain (on which Exchange is running) to Windows 2000 and connect to an existing forest.
- Upgrade existing servers and move them to a new Windows 2000 domain.

Do You Need to Upgrade All Servers?

Because upgrading takes time, labor, and money, you may decide not to upgrade the operating system on all your servers before you move to Exchange 2000. Consider the following scenarios:

- Existing server hardware may not be on the Windows 2000 hardware compatibility list. Typical examples are alpha-based servers and older servers based on proprietary processors. You may need to update some components to meet Windows 2000 requirements, but the upgrade options might not be strategic or cost-effective for your organization.
- You may be planning to retire an older generation of hardware during your deployment.
- You may be planning to consolidate on larger servers. To do this, you would need to move mailboxes anyway. Because you can move mailboxes directly to an Exchange 2000 server from any earlier version of Exchange, there is no reason to upgrade all existing Exchange *5.5* servers.
- You may want to begin running multiple smaller databases on the same hardware. This also requires moving mailboxes, so although you plan to re-use servers, you need to stage the upgrades with at least one extra leapfrog server. Although you plan to upgrade servers to Windows 2000, there is no need to do so while running Exchange 5.5; it may be easier and less risky to completely re-install each server after all mailboxes have been moved from the server.
- You may want to build a fresh environment and do so by using a leapfrog upgrade approach

In all of these situations, there is no need to upgrade to Windows 2000 Server in advance of your upgrade to Exchange 2000. In cases in which you plan to re-use hardware, you can reformat the hard disk and install Windows 2000 Server. If this is the case, you can skip the remainder of this section and read "Upgrading Exchange Mailboxes," later in this chapter.

an Domain and t to a Forest

This section assumes that an existing Windows NT 4.0 domain containing Exchange servers will be upgraded to Windows 2000 and joined to the existing forest as a child of an existing domain or as the root domain of a new domain tree. This section is an overview of specific planning issues that affect Exchange servers in a resource domain; you should also study the relevant Windows 2000 Server documentation before planning your domain upgrade.

This scenario is most likely to occur in an organization in which all the Exchange servers are installed in one or just a few resource domains. In many cases, some or all the Exchange servers will be backup domain controllers, and in some cases, the primary domain controller may even be an Exchange server.

To upgrade the resource domain, you start by upgrading the primary domain controller to Windows 2000 Server. This promotes the server to a domain controller and joins the domain to the existing forest. The upgraded primary domain controller is now also a domain controller in Active

Directory; it holds a replica of the schema and configuration naming contexts and the local domain naming context. The server does not automatically become a global catalog server, because by default, only the first domain controller in the forest is a global catalog server.

Next, you can upgrade the backup domain controllers in the domain. The steps to follow depend on whether or not you want the server to be a domain controller in the Windows 2000 domain.

At the end of the upgrade of a backup domain controller, Active Directory Installation Wizard runs automatically. If you want the server to be a Windows 2000 domain controller, then you follow the steps to join the existing domain. If you don't want the server to be a domain controller, you still step through Active Directory Installation Wizard and join the domain, but after that, you run Active Directory Installation Wizard again and demote the server to member status in the domain.

Design Issues

There are some important design issues to consider before you begin the upgrade process:

- How many domain controllers do you need in the upgraded domain?
- Will some or all of these domain controllers be global catalog servers?
- Will any Exchange 5.5 servers also be domain controllers?
- Will any Exchange 2000 servers also be domain controllers?

Much depends on whether the existing primary domain controller and any backup domain controllers are also Exchange servers. If they are not (all existing Exchange servers in the domain are member servers), the domain upgrade process is relatively straightforward: complete your design (which includes the number, location, and specific role of Windows 2000 domain controllers and specifications for required hardware upgrades), then upgrade the primary domain controller, and finally upgrade the backup domain controllers.

Exchange Running on Primary or Backup Domain Controllers

If your primary domain controller and backup domain controllers are Exchange servers, you may want Exchange to run on a domain controller after the upgrade.

If you don't want Exchange 2000 to run on a domain controller, then you must install new servers as Windows 2000 domain controllers. The following procedure represents the most economical upgrade path. (To simplify the discussion, it is assumed that the current primary domain controller and all backup domain controllers are Exchange servers.)

1. Install one of the new servers as a backup domain controller in the domain.
2. Promote this server to the primary domain controller and demote the existing primary domain controller to a backup domain controller.
3. Upgrade the new primary domain controller to Windows 2000, run Active Directory Installation Wizard, and then join the Active Directory forest.

4. Install the rest of the new servers in the domain as Windows 2000 domain controllers.
5. Upgrade all remaining Exchange servers to Windows 2000 and run Active Directory Installation Wizard to demote the server from a domain controller to a member server.

If an existing server doesn't meet your hardware requirements for Exchange, move all Exchange 5.5 mailboxes from the server, decommission the server, and remove the backup domain controller from the domain.

If you want your Exchange servers to be Windows 2000 domain controllers, you must upgrade the primary domain controller first, and then upgrade the backup domain controllers in any order. Before you do so, however, consider the following design issues:

- If you have a large number of Exchange servers, consider carefully whether the hardware can run both Exchange 2000 and manage Active Directory replication.
- You might need to upgrade your hardware by increasing the memory and processor capacity of your servers, especially if you plan to make any of the servers global catalog servers.
- There will be contention for the Lightweight Directory Access Protocol (LDAP) port. You need to change the default LDAP port (TCP port 389) for Exchange to a different port (because Active Directory binds to this port by the time the server has booted). You must update the configuration of any LDAP clients that connect to the Exchange directory. On Exchange 2000 servers that have Site Replication Service (SRS) enabled, you must change the LDAP port used by SRS.

Upgrade Servers and Migrate to a Forest

This scenario assumes the following:

- You want to use the in-place upgrade method to upgrade from Exchange 5.5 directly to Exchange 2000.

  You want to put the Exchange 2000 servers in a different domain.

Unless you plan to run an in-place Exchange upgrade, there isn't a strong reason to upgrade to Windows 2000 while running Exchange 5.5, especially if you are using a move mailbox method to move existing mailboxes to Exchange 2000 servers. It is easier to run a fresh installation on a server targeted as an Exchange 2000 server.

Only Windows 2000 member servers can be moved between domains. If Exchange 5.5 is running on member servers, then those servers can be upgraded and moved, or moved and upgraded. You can use NETDOM, a tool provided with the *Windows 2000 Server Resource Kit* that allows management of Windows 2000 domains and trust relationships from the command line, and Active Directory Migration Tool.

If Exchange 5.5 is running on backup domain controllers (and possibly the primary domain controller), those servers must be upgraded and demoted before they can be moved. A backup domain controller can't be upgraded if the primary domain controller has not already been upgraded. There are several approaches to migrating a domain:

- Set up a primary domain controller on a separate (non-Exchange) server, upgrade it, and make the domain the root domain of a new forest. You could join the existing forest, but this might cause problems if accounts from the domain have been cloned. The backup domain controllers can then be upgraded, demoted to members, and migrated.
- Each backup domain controller can be taken offline, promoted to a primary domain controller, upgraded, demoted to member status (in a workgroup), brought back online, and migrated to the target Windows 2000 domain.

You may also need to change the Exchange service account for all of the Exchange sites supported by servers in the domain from which you are migrating. If this domain is to be removed, then you must have a new service account in the Windows 2000 forest.

Upgrading Exchange Mailboxes

As illustrated in the deployment roadmap, you have three options during this phase:

- **In-place upgrade** This involves upgrading the server and databases on a single computer.
- **Move mailbox upgrade** This requires that a target Exchange 2000 server be installed in an existing Exchange 5.5 site. After that, you can move mailboxes by using Active Directory Users and Computers in Microsoft Management Console (MMC). As with the Exchange 5.5 move user scenario, single-instance message storage is preserved after you move mailboxes. This can save considerable disk space.
- **Leapfrog upgrade** This is essentially the same as moving mailboxes, except for the logistics. The goal is to reuse existing servers; a single Exchange 2000 sewer is added to the Exchange 5.5 site, then all the mailboxes on one Exchange 5.5 server are moved, and finally the Exchange 5.5 sewer is reformatted and installed as an Exchange 2000 server. You move mailboxes from another Exchange 5.5 server to this Exchange 2000 server. This process continues until you have moved all Exchange 5.5 mailboxes to Exchange 2000 servers.

The upgrade method you select depends on what you plan to do with your hardware.

- If you plan to deploy new hardware, then you will probably find it easiest to move mailboxes.
- If you need to upgrade hardware, or would like to reconfigure or reinstall, then a leapfrog upgrade may be appropriate.
- If your hardware is adequate in its current configuration, you may find in-place upgrades the easiest to manage. In-place upgrades are likely to be the most economical, especially in large distributed networks in which small sites support single servers.

To assist you in cases in which the path is not obvious, the following table lists the advantages and disadvantages of each method. In practice, you will probably need to investigate your own circumstances more carefully to estimate the costs involved. You can use the information in the following table as a general guideline.

Table 6.2 Upgrade method advantages and disadvantages

Method	Advantages	Disadvantages
In-place upgrade	Does not require extra hardware	Downtime
	The only option, if you don't have spare servers (especially at remote, single-server sites)	Risk associated with upgrading the database format
	Potentially the most economical option	Difficult recovery scenarios
Move mailbox upgrade	Minimal impact on users	Requires extra hardware; amount of extra hardware available constrains rate
	Easy recovery	Slow; large servers could take days or weeks
	The only option, if you are replacing hardware	Complicated logistics, especially for remote locations
	The most efficient option, if you are changing to multiple databases per server	
	Fresh start for server configuration, disk partitions, and databases	
Leapfrog upgrade	Minimal impact on users	Requires some extra hardware (biggest impact on small sites)
	The only option, with certain types of hardware upgrade or reconfiguration	Complicated logistics
	Fresh start	

Note An in-place upgrade involves an update of the database format from the ESE97 format used by Exchange 5.5 to the ESE98 format used by Exchange 2000. Upgrades from earlier versions of Exchange probably involve an upgrade to ESE97 format before the upgrade to Exchange 2000.

The information in the following paragraphs can help you measure the risk and benefit involved with databases:

- **Risk** Previous Exchange database upgrades have involved a complete pass through the database, updating all tables and items to the new format. The upgrade from ESE97 to ESE98 does not do this; the upgrade touches only a few key areas of the database. After the upgrade, items (rows in tables) are updated dynamically as they are touched during normal operation of Exchange. This approach results in significantly better upgrade performance; it should not cause problems after the upgrade. It does, however, mean that an upgrade of a database cannot be tested in the same way that it could for previous upgrades. Future upgrades will probably follow the same pattern. While this is identified as a risk, it is an extremely low risk.
- **Benefit** There may be concrete benefits in making a fresh start. An example relates to the space tree in an ESE database. The space tree is a b-tree that tracks free space in the database. In an ESE97 database, the space tree itself is never shrunk; it only grows, and in large databases the growth of the space tree can result in performance problems. There are hotfixes available for these problems. ESE98 fixes this, and will shrink and optimize space trees. However, the optimization will take some time, and there could be a period immediately after upgrade in which performance degrades significantly.

Coexistence

After you introduce the first Exchange 2000 server to an Exchange 5.5 organization through an upgrade or a fresh installation, the Exchange organization is in ***mixed*** mode. Just as Windows 2000 has mixed mode, the organization must stay in mixed mode until all of the Exchange 5.5 servers have been upgraded or decommissioned.

Operating in mixed mode imposes certain restrictions, which affect administrators more than users. As soon as a user's server is upgraded to Exchange 2000, the user gets access to the full range of new features. The restrictions affect how the administrator manages the upgraded server.

In mixed mode, routing groups must correspond to administrative groups, except when all the servers in an administrative group have been upgraded to Exchange 2000. Then you can create additional routing groups in the administrative group, but they can't cross administrative group boundaries. After the organization has been switched to native mode, routing groups can be defined independently of administrative groups.

This restriction must remain in place in mixed mode, because there are two versions of the Exchange directory: one maintained by the Exchange 5.5 directory service and one in Active Directory. For message routing and related operations to work, the two versions must correspond, at least in terms of overall topology, and the version of the directory maintained in Active Directory must adhere to the rules for an Exchange 5.5 directory.

In addition, you cannot move Exchange 5.5 servers between Exchange sites. In mixed mode, an Exchange 5.5 site corresponds with an Exchange 2000 administrative group, which must also correspond with an Exchange 2000 routing group. If this relationship is broken, reliable message transfer will no longer occur.

Configuration Connection Agreements and Site Replication Service

The components that maintain two corresponding versions of the Exchange directory during mixed mode operation are ADC and SRS.

You install and configure an instance of the SRS on the first Exchange 2000 server installed in an Exchange 5.5 site, on the first server in a site to be upgraded, or on a bridgehead server. The SRS looks and behaves for the existing Exchange 5.5 servers like the Exchange 5.5 directory service—it participates in directory replication the same way as any Exchange 5.5 server. SRS provides the mechanism for transparently turning the Active Directory version of the organization's Exchange directory into an Exchange 5.5 site.

SRS allows the Exchange 2000 server to more closely resemble an Exchange 5.5 server. It includes the same Dir.edb file as the one used on all Exchange 5.5 servers for directory services. It also uses the same configuration connection agreements that are created when you install Exchange 2000. The Exchange 2000 Setup Wizard analyzes your Exchange organization and builds the configuration connection agreements necessary to replicate Exchange-specific configuration information between Exchange 5.5 and Active Directory. Configuration connection agreements are used to replicate Exchange-specific configuration information between Active Directory and the SRS.

The SRS informs the Exchange 5.5 servers about the entire Exchange 2000 configuration. This way, the Exchange 2000 servers that host the SRS can perform the Exchange 5.5 configuration replication. In addition, the SRS is displayed in the Exchange 5.5 Administrator program as the Exchange 5.5 Directory Service. For the Exchange 2000 servers that are not hosting an SRS, no Directory Service is listed.

Whereas SRS provides a connection to the Exchange 5.5 directory, ADC provides a connection to Active Directory. When the first Exchange 2000 server is installed, a special connection agreement called a configuration connection agreement is created automatically. The configuration connection agreement replicates the topology of the Exchange 5.5 organization to Active Directory, which is maintained in the configuration naming context. The configuration connection agreement also replicates changes made to the Active Directory version of the directory to the SRS.

Figure 6.1 shows an Exchange 5.5 organization comprising two sites with a directory replication connector between server 3 and server 4, and Windows 2000 Server installed on server 2.

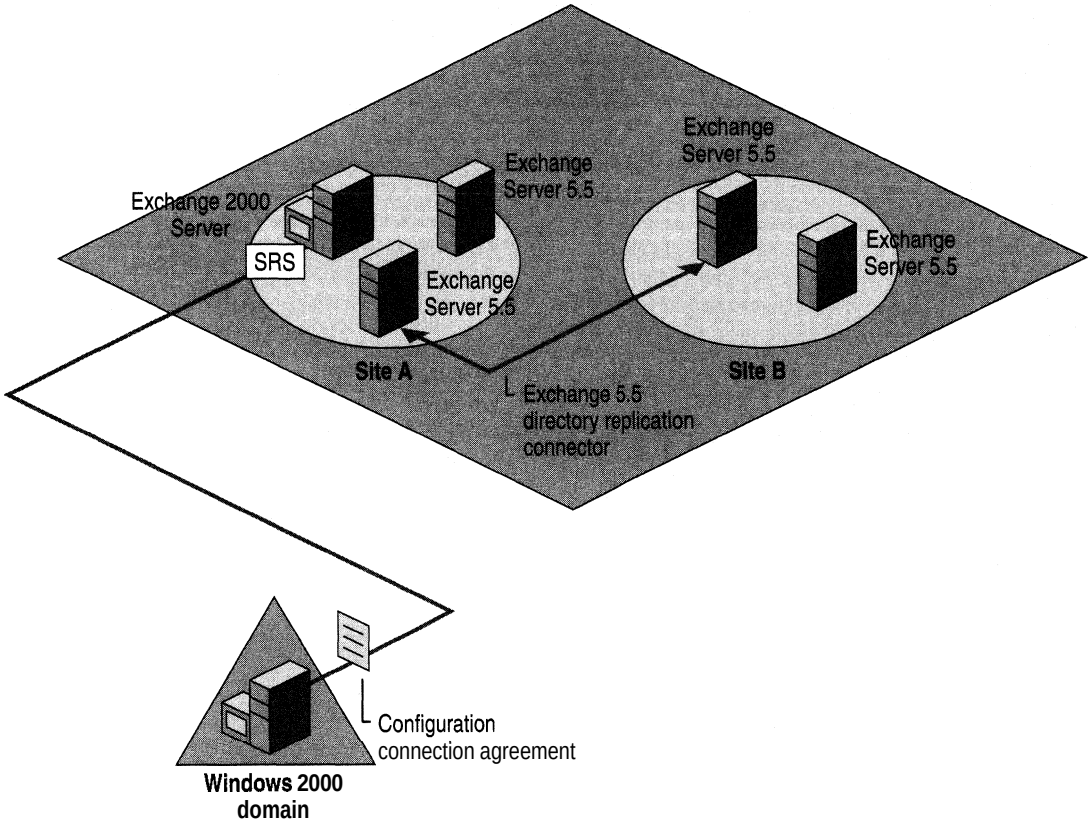


Figure 6.1 SRS is enabled on an Exchange 2000 sewer

In Figure 6.2, server 1 is upgraded to Exchange 2000. The SRS is also installed, but it remains disabled because SRS is already functioning on server 2.

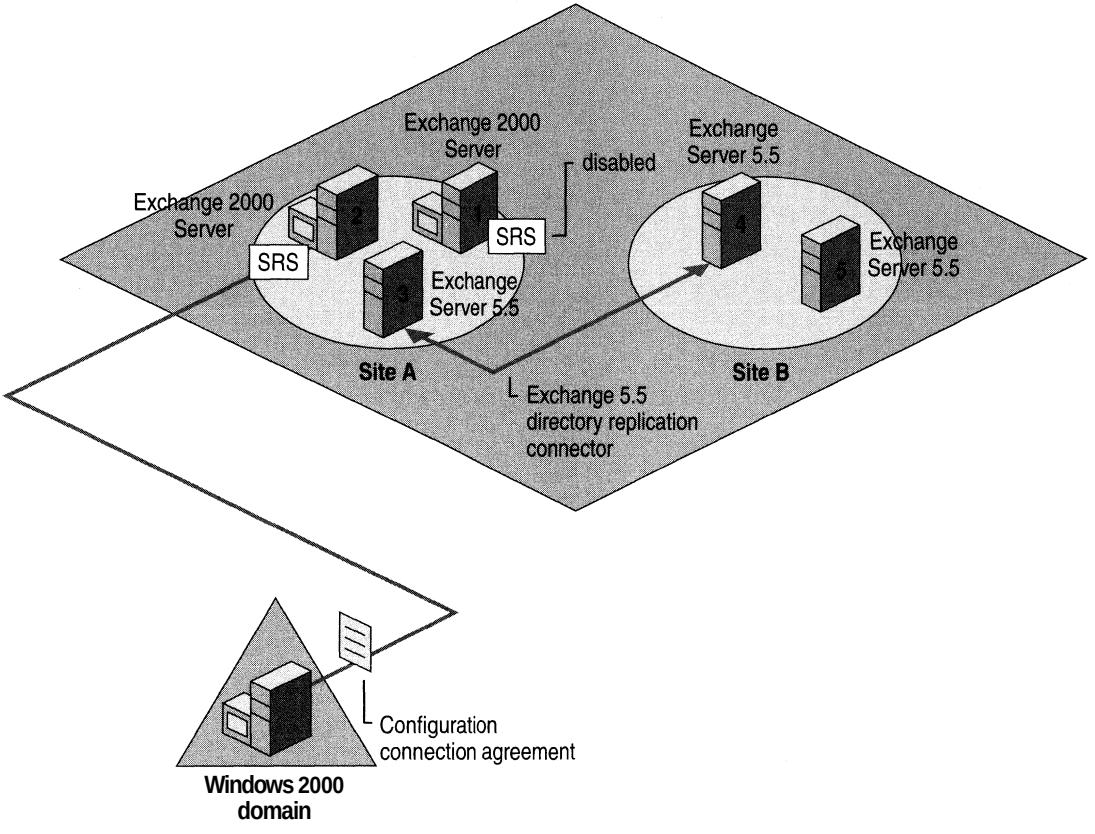


Figure 6.2 SRS is disabled on second Exchange 2000 server

Without another Exchange 5.5 server in the site or SRS on server **3,** server 3 cannot participate in Exchange 5.5 intra-site directory replication; all configuration information flows using the existing configuration connection agreement.

Figure 6.3 shows what happens when you upgrade server 3, the bridgehead server. After you upgrade server 3, SRS on server 3 allows replication between site A and site B to continue to work.

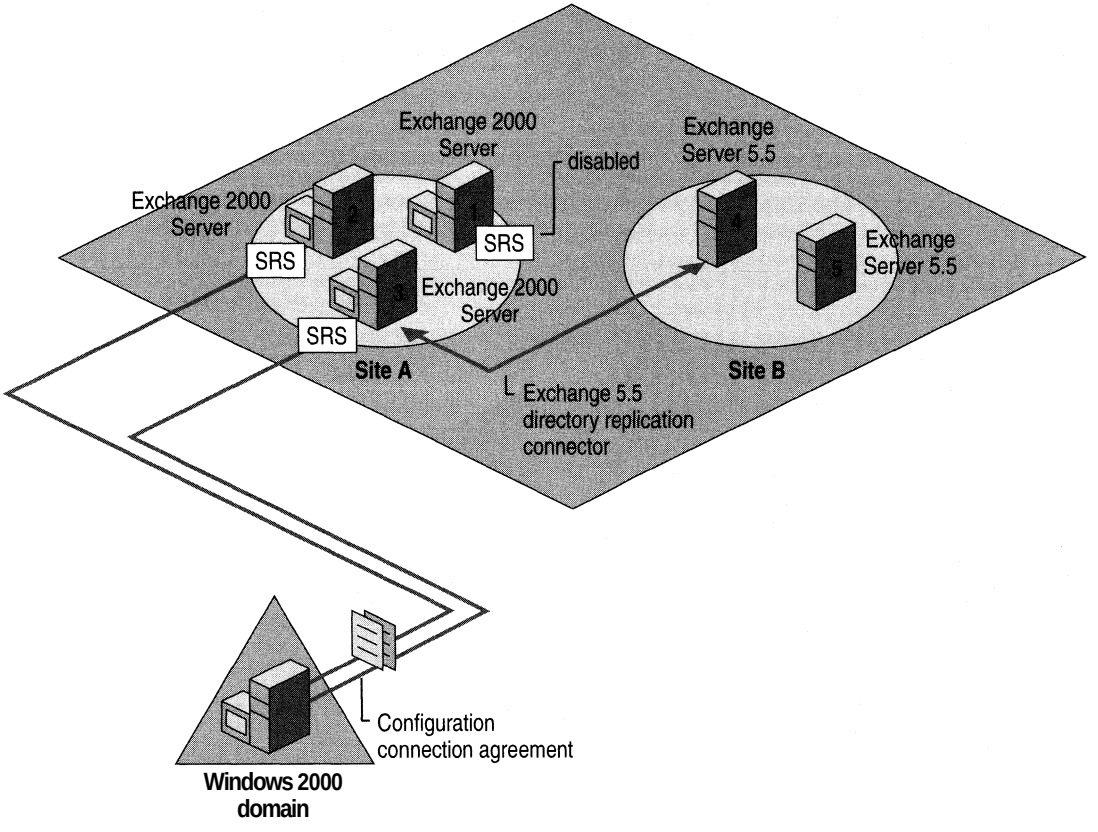


Figure 6.3 SRS is enabled on the directory replication bridgehead server

The second configuration connection agreement isn't necessary, but it makes the overall coordination process easier. A special Knowledge Consistency Checker that is installed with SRS, called Super Knowledge Consistency Checker, dynamically configures multiple configuration connection agreements to establish the most efficient replication. The implications from this discussion include:

- If the first Exchange 2000 server in the site is also the directory replication bridgehead server, then you will have only one enabled SRS and one configuration connection agreement. This simplifies the overall configuration.
- If you already have user connection agreements in the site, they must be moved before you upgrade or decommission the server on which they are located. If you move them to the SRS, you will not need to move them multiple times.

- Because SRS is enabled only on the first Exchange 2000 server and the upgraded bridgehead servers, the number of servers participating in Exchange 5.5 intra-site replication decreases as you continue to upgrade. Large sites benefit immediately from the reduced network replication traffic.

The order in which you upgrade servers on a site is typically determined by business priorities, such as the need to retire old hardware, consolidate mailboxes on fewer servers, or deliver new capabilities to specific groups of users.

Upgrading multiple sites poses additional challenges when you install a new server or upgrade an Exchange 2000 server on the second site, shown in figure 6.4. When server 5 is upgraded, SRS is enabled and a configuration connection agreement is created between the SRS and Active Directory. This configuration connection agreement is now responsible for replicating information about site B to and from Active Directory. The configuration changes are all handled by Super Knowledge Consistency Checker.

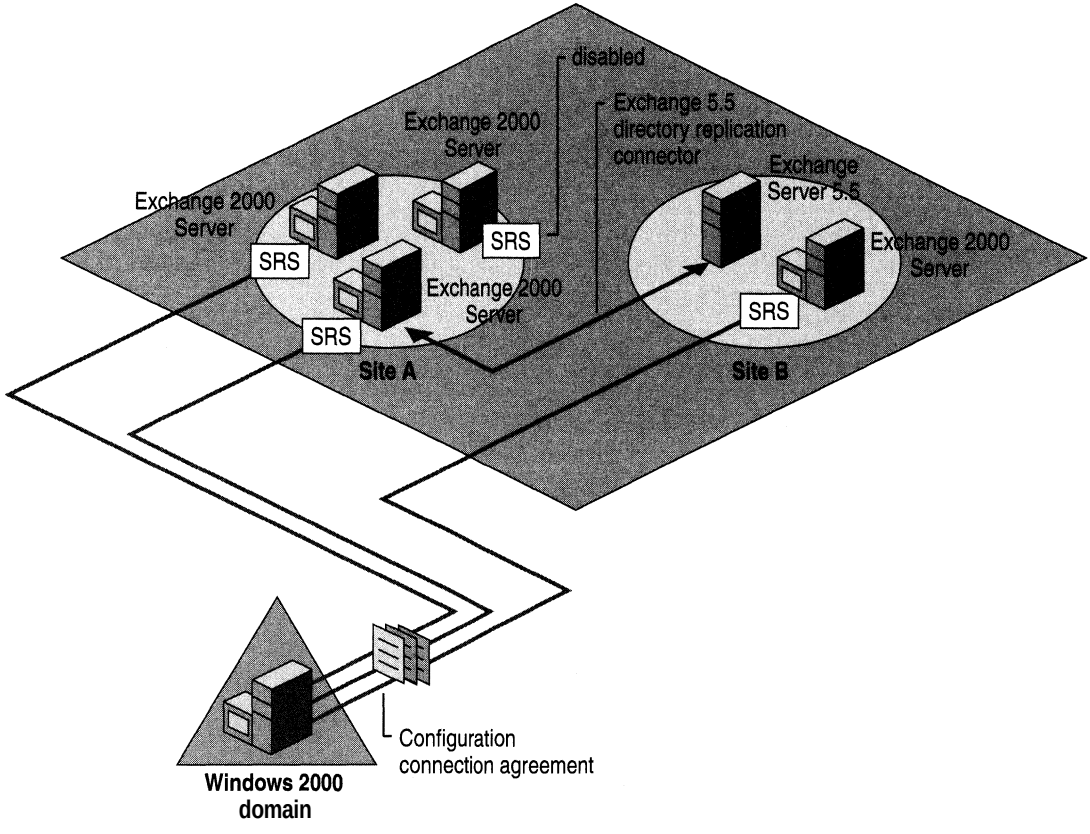


Figure 6.4 SRS across multiple sites

All three configuration connection agreements exist and continue to work

In organizations with many sites, Super Knowledge Consistency Checker performs a central function called *arbitration.*

In a new example, assume that site A is upgraded first. An SRS and a configuration connection agreement are installed on the first server. The configuration connection agreement replicates the Exchange 2000 part of site A inbound to SRS, and the Exchange 5.5 part of all six sites outbound from SRS to Active Directory.

Next, site C is upgraded. An SRS and a configuration connection agreement are installed on the first server. The configuration connection agreement for site C replicates all information inbound to and outbound from the SRS for site C. The process of arbitration determines which SRS is responsible for replicating the configuration information for site B to Active Directory. Super Knowledge Consistency Checker performs arbitration.

Note There is no user interface to view or manipulate the results of arbitration.

Arbitration is based on site name. When Super Knowledge Consistency Checker needs to decide which of two sites should take responsibility for a third site, it selects the site with the closest alphabetical name.

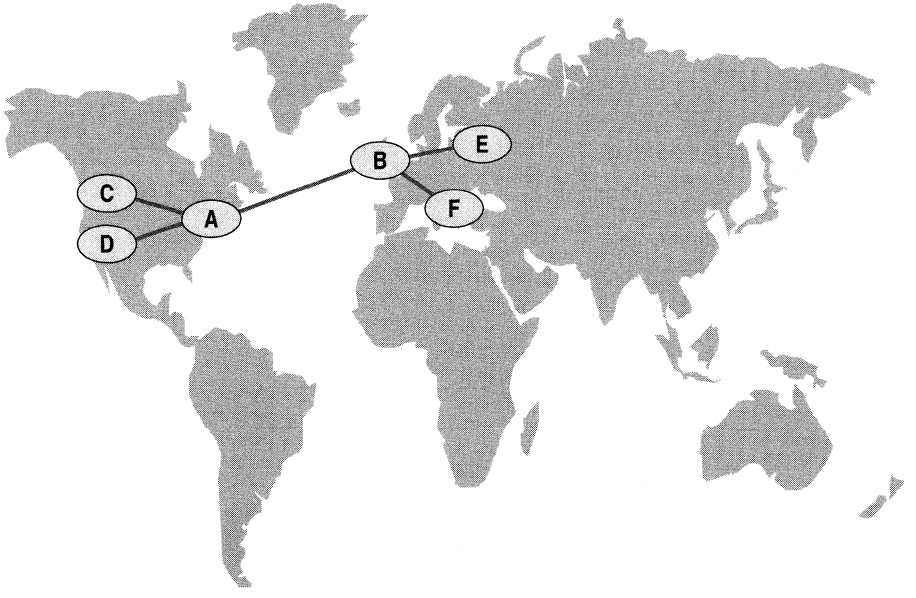


Figure 6.5 Upgrading an Exchange 5.5 organization with distributed sites

Arbitration has ramifications for distributed multi-site Exchange 5.5 networks. Consider a typical site topology for a trans-Atlantic organization, as illustrated in Figure 6.5. Sites A and B act as hub sites, both for message traffic and for directory replication. If the first site upgraded with an Exchange 2000 server is site C, site C is responsible for replicating all the other sites to Active Directory. If the next site upgraded with Exchange 2000 is site B, site B is responsible for replicating itself, but site C continues to replicate sites E and F. Because site A separates site C from sites E and F, this isn't the most efficient replication topology.

With simpler network designs, this problem can be overcome by carefully choosing which site to start with. If you have a single hub site, then you can install the first Exchange 2000 server into the hub site and the replication topology will be no worse than it is at present. You may be able to reduce network replication traffic by starting with local sites with local Windows 2000 domain controllers.

Distribution Lists and Permissions

The transition from an Exchange-specific security model in Exchange 5.5 to the Windows 2000 security model in Exchange 2000 poses special challenges.

In Exchange 5.5, you can assign permissions to any type of recipient in the global address list. Some administrators use distribution lists widely to build systems for access permissions to public folders.

With Exchange 2000, you can assign permissions only to a Windows 2000 user or security group. When Exchange 2000 upgrades the public folder store and mailbox store, it must be able to replace Exchange 5.5 distribution lists that appear in ACLs with Windows 2000 security groups.

However, some additional complications arise. When upgrading to Exchange 2000 you must consider the following facts:

- ADC replicates Exchange 5.5 distribution lists to Active Directory as distribution groups, which are not security groups.
- Exchange 2000 needs to evaluate group membership when applying permissions, but only one type of Windows 2000 group has its membership replicated to all global catalog servers: the universal group.
- Universal security groups can be created only in native-mode Windows 2000 domains.

Thus, during the upgrade process, Windows 2000 universal security groups must replace Exchange 5.5 distribution lists. For this to be possible, there must be at least one native-mode Windows 2000 domain in the forest in which the security groups can be created. The entire forest doesn't need to be in native mode; you can create a special native-mode child domain specifically for the purpose of holding universal security groups. This domain is referred to as the transition domain.

To facilitate your upgrade to Exchange 2000, it is recommended that Windows 2000 domains be converted to native mode as early as possible. In addition, cloning of users requires that the target domain be in native mode (the sIDHistory attribute is only available in native mode).

Services That Cannot Be Upgraded

There are several Exchange 5.5 services that can't be upgraded immediately to Exchange 2000. They include:

- Exchange 5.5 SNA Distribution System (SNADS) and Professional Office System (PROFS) connectors. Microsoft does not produce Exchange 2000 versions of these connectors. Each server running one of these connectors must continue to run Exchange 5.5 until the connector is no longer required.
- Third-party connectors, such as FAX connectors, for Exchange 2000 that are not available when Exchange 2000 ships. You might need to wait for the third-party vendor to produce an Exchange 2000 version of the connector before you upgrade.
- Antivirus software. As with third-party connectors, you might need to wait until upgraded versions are available before you upgrade antivirus software.

Scenarios

The boundaries of an Exchange 2000 organization are determined by the boundaries of the Active Directory forest, which are also the boundaries of the replicated schema and configuration naming contexts.

If you need to deploy Exchange 2000 in an environment in which multiple forests exist, you can't deploy a single Exchange 2000 organization. You have two choices:

- Deploy separate Exchange organizations in each of the existing (or planned) forests and synchronize the directories.
- Deploy Exchange in its own (new) forest, which provides only mailbox services.

Separate Exchange Organizations

If you plan separate Exchange organizations, two questions arise: How do you set them up and what features do you lose?

Setting up multiple Exchange organizations requires synchronization. Because recipients in other forests will not be security principles in the local forest, they must be maintained as contacts (the Active Directory equivalent of Exchange 5.5 custom recipients).

You can use several tools to synchronize between forests, including ADC. When you create connection agreements to synchronize between forests, you can specify that the connection agreement is an inter-organization agreement. This allows ADC to synchronize with a directory that is not part of the same organization. You can also specify that users were created as contacts in the connection agreement configuration.

What you lose with multiple forests is similar to what you lose with multiple Exchange 5.5 organizations:

- Access to public folders
- Access to calendar free and busy information
- Access to other users' appointment books

Additional losses that relate to new features in Exchange 2000 include conferencing services and instant messaging. There are also other restrictions, both in terms of delegated permissions and access to applications, but they depend on how you use Exchange in your organization.

Though closer integration with the Web in Exchange 2000 offers some solutions, there is a new restriction when compared with Exchange 5.5. You cannot assign permissions to Active Directory contacts. They are not security principles (if they were, they would be users). This is different from Exchange 5.5, which implemented its own security model. So although closer Web integration can solve some of the visibility restrictions, it is likely to do so in ways that can compromise security.

Exchange 2000 in Its Own Forest

In the same way that Exchange 5.5 can be implemented in resource domains that trust separate Windows NT 4.0 account domains, Exchange 2000 can be implemented in its own Windows 2000 resource forest that trusts account domains in other Windows 2000 forests. This scenario is illustrated in Figure 6.6.

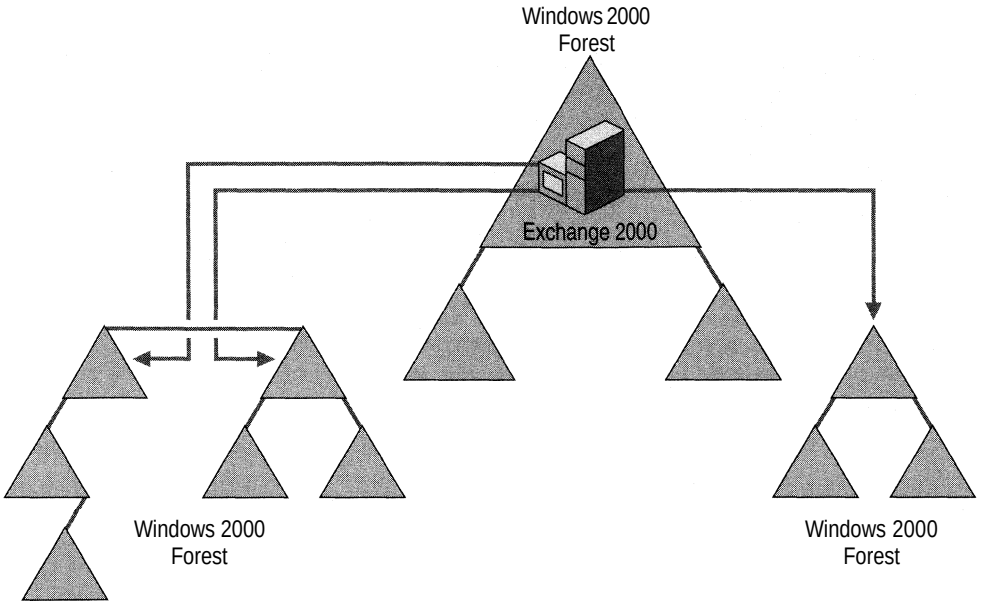


Figure 6.6 Exchange 2000 in a separate forest

The idea behind this scenario is that users log on to their home forest and use credentials from that forest to log on to Exchange. There are two requirements for this to work:

- The user's account must have user permissions to the Exchange 2000 mailbox.
- A down-level trust is required from the domain containing the Exchange servers to every domain in every forest that contains users with Exchange mailboxes.

If your Exchange 5.5 servers are in one or more resource domains, this design is relatively easy to implement. You complete the following steps:

1. Upgrade the Windows NT 4.0 resource domain to a Windows 2000 domain in its own new forest. If you have multiple resource domains, you can upgrade the rest and join the forest, or migrate servers across forests.
2. Install ADC and set up a connection agreement to create users in the new forest. By default, ADC creates disabled user accounts. They are mailbox-enabled users, and the primary Windows NT Account from the Exchange 5.5 directory is added to the mailbox security descriptor attribute.
3. Upgrade Exchange servers, or install new Exchange 2000 servers and move mailboxes. The upgrade process doesn't touch the mailbox security descriptor, so as soon as a user's mailbox is on an Exchange 2000 server, the user can log on using his or her existing credentials (in Windows NT 4.0 or Windows 2000).

This result could be achieved by other methods, but all methods must assign mailbox permissions to the user's existing security credentials. This approach is advantageous because all users are in a single Exchange organization. Visibility of public folders, free and busy information, and appointment books is assured, and integrated with a fully functioning security model.

The disadvantage of this method is that there is another forest to manage with another set of accounts. ADC can create this environment easily if you start from an Exchange 5.5 organization, but after you are running Exchange 2000 throughout your organization, you will need procedures and processes to synchronize the disabled user accounts with users' primary accounts in the forests they log on to.

To make this manageable, you may need a permanent directory synchronization infrastructure based on a flexible tool like Microsoft Metadirectory Services (MMS) version 2.2, which allows filtering and mapping of attributes. At the time this chapter was written, this version of MMS was not yet available.

Multiple Exchange 5.5 Organizations

To merge two or more Exchange 5.5 organizations, you can use the Move Server Wizard.

Rather than waiting until you upgrade to Exchange 2000, you should merge organizations while still on Exchange 5.5. The primary reason for this recommendation is that there is no way to merge two Exchange 2000 organizations. You can migrate user accounts and Windows servers between forests, but you can't migrate Exchange databases.

It is possible to use the Exchange 5.5 Move Server Wizard to migrate Exchange 5.5 servers into a mixed-mode Exchange 2000 organization. However, you can do so only under the following conditions: the server you are moving must exist in a site with Exchange 5.5 only, and there must be at least one Exchange 5.5 server already running in the target site.

PART 3

Prototyping Exchange 2000

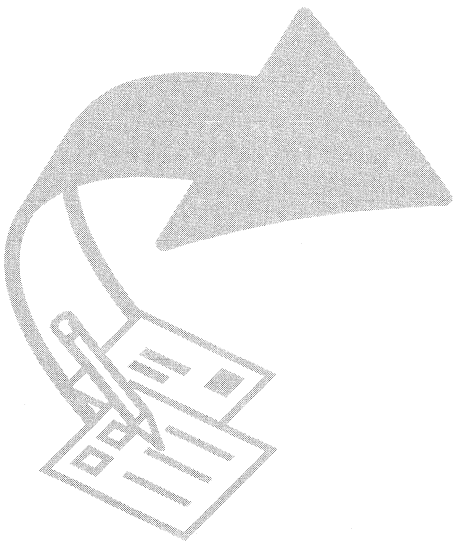
In This Part

Setting Up a Test Environment

Piloting Exchange 2000

Preparing a New Environment

Preparing an Existing Environment



CHAPTER 7

Setting Up a Test Environment

Bob Hunt, Managing Consultant, E-Sync Networks, Inc.
Paul Sebben, Managing Consultant, E-Sync Networks, Inc.

Setting up a test environment to evaluate new technologies is a common practice for successful deployments. When new, complex technology is business critical and affects many users, it is prudent to do a thorough product evaluation. When you perform a complete product evaluation and a scenario-based test plan, you will gain knowledge of a product and how it will react in different situations. This is the best insurance plan your company can have for an enterprise-wide implementation.

The tests need to put the product through most of the scenarios relevant to your company's implementation of Microsoft Exchange 2000. This chapter provides information about setting up test environments and putting together thorough test plans for Exchange 2000. This chapter also cites specific features that should be tested in various scenarios relevant to an enterprise-wide implementation of Exchange 2000.

In This Chapter

Why Testing Is Important

Planning the Test Lab

Gathering Requirements

Developing a Test Plan

Testing Your Scenarios

Why Testing is Important

With any enterprise software deployment, the importance of testing and establishing a relevant testing environment cannot be stressed enough. You should not worry that a specific feature does not function as advertised, but rather how a specific feature affects the enterprise. At the highest level, a test environment allows you to install and configure a new product prior to its implementation. Many enterprise-wide application deployments are complex in size and scope; the depth of which may not be understood until actual implementation. However, gaining in-depth knowledge of a product prior to implementation reduces risk because each phase of implementation has been tested in a similar environment.

For example, there are various connectors available in Exchange for linking locations across a company's WAN. Knowledge of the connector network overhead provides the information that you need to select connectors, based on available bandwidth between locations.

When you install software in specific test scenarios that are similar in scope to the actual enterprise-wide implementation, you can perform one or more trials. This reduces the likelihood of potential failures in an actual network environment. Your test environment does not have to completely duplicate the company's network infrastructure, but it should include the major architectural components necessary for validating the majority of the implementation scenarios. If the test lab closely mirrors the production environment, tests performed in the lab can set expectations for the behavior of the application's production environment.

Planning the Test Lab

The goal of the test lab is to establish a model of the pre-implementation production environment. The first step is to identify what comprises a small-scale model of your network infrastructure. Keep in mind that simultaneous testing of several production environment scenarios is desirable, but is not necessarily a requirement when planning for the test lab.

Realistic budgeting is essential. Although the acquisition and construction of a test lab can be expensive, it is important to avoid using some or all of the lab equipment for the production environment. It is important to understand that the need for a test lab does not go away after the implementation is complete. Convenient and easy to use disk imaging software enables companies to build lasting and substantial test lab environments that can be used in their enterprise application deployments for many years. For example, companies must have the ability to assess the impact of new service packs in a test environment. New applications, including third-party add-ons, can be installed in the test environment for evaluation and testing purposes. Another use for the test lab environment is as a training center for new employees or administrators. Test labs can also be used for application fine-tuning, network performance and analysis, backup and restore, and disaster recovery.

Install and maintain the lab as if it were a production environment. Investing in a test lab allows your support staff to test administration, support, maintenance, and troubleshooting techniques without affecting the production environment. The lab itself should always be ready to test new products or version upgrades, patches, and service packs. You can appoint a lab manager or coordinator to ensure control over the lab conditions. This person can be the primary contact for the lab configuration, and should have signature authority on changes to the environment.

A system of change control is essential to keep the lab in a continuously stable and known state. The lab manager can ensure that the testing efforts of multiple groups do not interfere with each other. The lab manager is also in charge of the process to image and restore computers to their original state. Include a step for verifying changes in the test environment.

Office Space

When you consider space for your test lab, select a location that can be kept clean, cool, and undisturbed. Tables, chairs, and other office furniture are necessary so the lab environment is usable and dependable: this furniture should be dedicated to the lab. Find a secure location where test scenarios can be left running without disturbance from other users, and where hardware is secure.

Hardware

Approach choosing hardware for the test lab in one of two ways. First, you can ensure that the hardware you choose is identical to the hardware used for production deployment. The advantages of this strategy are that it allows you to perform server scalability and performance testing, and to use lab hardware as backup for the production environment. A disadvantage is high cost. Having enterprise-class hardware sitting idle may not fit the corporate information technology budget.

The second option is to choose hardware that is less expensive than the enterprise-class hardware that you intend to use for the deployment. The primary advantage of this option is the hardware for the test lab can be obtained at a relatively low cost. Disadvantages to this approach include, not gaining familiarity with the hardware to be used during the deployment, and not being able to perform hardware scalability and performance tests.

The most important factor to keep in mind is that the lab must contain the major components of the company's production infrastructure for relevant software testing. If resources are not available for the test lab, you can still perform specific software implementation scenarios on less capable computers. Early acquisition of a production-class server allows for scalability testing on the chosen hardware prior to implementation. The company can then verify performance, mailbox count, server size, and other potential issues that need to be addressed.

Microsoft TechNet and the Exchange Web site at http://www.microsoft.com/exchange contain documentation on scalability, users per server, and other benchmarks. Some server hardware vendors may have information on Exchange server sizing on their company Web sites. All of these resources will give you information on the size of a server to be purchased.

Backup hardware for the test servers is a valuable component of any test lab. In addition to testing backup software and hardware, the test lab can be used as a mailbox recovery location if a mailbox needs to be restored. The company standard for backup software and hardware needs to be tested together and put through a number of backup and recovery scenarios so the organization is prepared for an unexpected loss of data.

Dedicate client workstations to the test lab environment using standard company software, and each type of messaging client in the organization. Include all versions of Microsoft Outlook supported by the information technology group.

Networking

As discussed earlier, the test lab's network environment must include the major components of the company's actual production network for relevant test scenarios. While it is preferred that identical network hardware that is in use in the production environment is procured for the test lab, duplicate functionality is most important.

It is important to keep the test environment completely separate from the production environment and on its own network. This will ensure that testing activity does not affect the production LAN or messaging environment. The only shared component from the network infrastructure (between the testing and production environments) should be Internet connectivity, unless an entirely separate connection is available.

It is important to set up a separate mail-testing domain for the testing environment that is configured to send and receive mail from the Internet. This is because Internet mail capability is a requirement for company e-mail systems. For example, if an e-mail domain name is microsoft.com, enabling a domain named test.microsoft.com requires making only a few Domain Name System (DNS) entries and any necessary changes to the firewall. It may be necessary to set up a cache-only or secondary DNS server for the testing environment for the separate root domain controller that holds the test environment's Active Directory.

An additional option to consider is how the production network will look in the future. Equip the test lab network with network upgrade hardware prior to production implementation so the impact of the upgrade can be measured.

Software

As mentioned previously, try to duplicate the company's network environment to create a model of the production environment. Because it is important that this model be separate from the production environment, the software resources located on the production network should be duplicated in the test lab. This includes domain controllers, Dynamic Host Configuration Protocol (DHCP) and Windows Internet Name Service (WINS) servers, any Microsoft Windows NT 4.0 servers with Microsoft Exchange Server 5.5, a computer running Microsoft Windows 2000 with a copy of the company's root Active Directory, and so on.

Have available in the test lab any third-party applications that might affect the production Exchange environment. This includes fax gateways, public folder add-ons, and backup software and agents. In addition, any e-mail systems that interact with Exchange in the production environment should be duplicated in the test lab. E-mail clients that are used throughout the organization also need to be made available in the test lab. These can include Outlook, Internet Message Access Protocol version 4 (IMAP4), Post Office Protocol version 3 (POP3), Microsoft Outlook Web Access, and Network News Transfer Protocol (NNTP) e-mail clients.

Gathering Requirements

Before testing begins, you should have a full understanding of the proposed messaging environment. A good place to start is to identify the shortfalls of the current environment ranging from significant issues, such as switching between different messaging systems, to specific drawbacks of the current environment. An example is the need to change the Exchange-organization name due to a company name change or an acquisition. If you are currently planning your test environment, it is likely that the requirements gathering phase for the Exchange 2000 messaging environment is complete.

Developing a Test Plan

A test or migration plan is essential for your company's successful Exchange environment implementation. It is important at this stage to understand all aspects of the product intended for the future production environment, and that this process of gathering information should end before you draft a test plan. All features, protocols, and connectors identified in the requirements gathering phase should be included in your test plan, so all architectural components of the production environment are tested in the lab.

Identify Risks and Contingencies

Because all activities related to the production deployment cannot be tested in a lab, a staff with a working knowledge of the product is insurance against potential disasters. While some problems may occur, thorough product knowledge prepares a team for unexpected issues.

Network Effects

An unavoidable risk during implementation is how changes to the messaging environment will affect the network infrastructure. An environment is only as strong as the infrastructure it runs on; unfortunately, you cannot test or know how the messaging infrastructure is going to impact the network infrastructure prior to the implementation.

As a contingency plan, you should test as many network-related items as can be duplicated in a test lab while considering the impact of adding the messaging components to the environment. A simulated test WAN on your network that measures messaging performance is valuable, but is not altogether conclusive without further information, such as the amount of messaging and replication traffic generated and sent to each site.

Implementation teams should gather network performance information from the company's WAN to assess available bandwidth, rather than identify the total link bandwidth. You should analyze the available bandwidth between each site during business and non-business hours to determine WAN traffic patterns. You should also compile a detailed protocol-level analysis of the company's WAN to determine the amount of e-mail, replication, standard network, and other WAN traffic. After you have gathered this information, you can test messaging components with a simulation of the anticipated bandwidth.

Administrative Access

Another risk during an Exchange 2000 deployment involves changes to administrative access. Administrative groups in Exchange 2000 allow specific control over administrative rights, but like many other security features, they must be thoroughly planned and tested. Doing this in the test environment allows information technology personnel to understand the power and flexibility of administrative groups. For example, if a company has two branch offices and within one are multiple divisions that share the LAN with the second branch office, you can set up separate administrative groups for each branch office. This allows users to modify only their own branch servers.

SMTP Connector

A critical detail in this process is the vital and potentially challenging Simple Mail Transfer Protocol (SMTP) Connector upgrade. This component is a core feature of Windows 2000, as opposed to a Microsoft Exchange 5.x add-on. It can be helpful to perform a test server in-place upgrade to understand the impact of the changes before production deployment.

A final consideration when upgrading is that when Exchange 5.5 public folders are replicated or upgraded to Exchange 2000, any associated distribution groups are converted to universal security groups. Even in native-mode Exchange 2000 organizations, when an administrator or a client changes the permissions on a public folder, they can create a distribution list and set permissions to allow members of that list to access the public folder. When you set permissions on a public folder, all associated distribution groups convert to universal security groups, and a native Windows 2000 domain must host them. Remember this restriction when assessing the risks of upgrading public folders.

Replication Latency

When you are performing potentially risky tasks. be aware that it might take time for your changes to replicate and become available. For example, if you set permissions on a domain controller, it may take a short period of time before the permission settings are replicated to other domain controllers in the same domain and Windows 2000 site. However, it is possible for Exchange services to request permission-related data from a domain controller without the most current replication information. Thus, when you make configuration changes such as adding new users and servers or moving mailboxes, you may experience a lag time before the settings are replicated to other domain controllers in the same domain and site. Site-to-site replication depends on your schedule for it. If you schedule replication every eight hours, changes in one site will take up to eight hours to replicate to all other sites.

Training

Training users is essential throughout enterprise deployments. Users who are unfamiliar with new software and functionality are a significant burden on the help desk and other support staff. This situation is easily avoided by establishing a training and ongoing support program for the users prior to production deployment. In addition, telling users where to find product documentation and general information is essential to the success of the project.

Third-Party Connectors

A final area of concern is third-party connectors, because there are many add-ons for Exchange 5.5 and earlier versions of Exchange that organizations use for core functionality. Many of these products are not compatible with Exchange 2000 and require thorough evaluation before upgrading. You can avoid this issue by keeping one or more Exchange 5.5 (or earlier) servers that run third-party connectors until new versions of connectors become available. After new versions become available, information technology personnel can identify migration scenarios to move this functionality to a new server.

Client Software Test Plan

Each client protocol in use must be included in the test plan to ensure compatibility. Some of the testing that should occur includes:

- E-mail delivery between Exchange 2000 and Exchange 5.5 servers in a different site.
- E-mail delivery between different e-mail systems over older connectors.
- Outlook offline and online client configurations connecting to global catalog servers to get directory information.

Include Outlook Web Access in the client test plan. Two different interfaces exist to provide access for browsers:

- Interfaces for earlier browser versions that use frames (Internet Explorer 3.02a. Internet Explorer 4.x, Netscape Navigator 3, and Netscape Navigator 4). Java is not included and there are fewer frames, resulting in better performance than Exchange 5.5 Outlook Web Access.
- An interface for Internet Explorer 5.0, which has a similar interface to Outlook.

You can test Outlook Web Access with different browsers and determine whether an upgrade to Internet Explorer 5.0 is required.

MAPI-based clients, such as Outlook 2000, Outlook 97, and Outlook 98 must be tested to ensure that you are providing the same level of support after Exchange 2000 is deployed. Testing issues include:

- **Basic messaging verification** Verify that users can send and receive e-mail while online and offline, read mail using preview pane, open mail, reply to a sender and to all recipients in the **To** and ***Cc*** boxes, forward mail, read various types of message content (attachments, embedded messages, HTML, Rich Text Format [RTF], plain text), and read high and low priority mail.
- **Name resolution** Verify that users can resolve other user names and distribution lists in the **To, Cc,** and **Bcc** boxes.
- **Copy and move** Verify that users can copy and move e-mail messages to and from personal folders (.pst), public folders, and that they can drag and drop messages in mailbox and public folders.
- **Folder operations** Verify that users can delete, rename, copy and move messages, set permissions, apply custom views, get properties, create subfolders, copy folder design, and work with folders offline.
- **Import and export** Verify that users can import and export .pst files.
- **Address book** Verify that users can view user properties, add users to their Personal Address Book, download Offline Address Book, and check user properties.
- **Rules** Verify that users can create, enable, or disable a rule.
- **Out of office** Verify that users can enable out-of-office notification.
- **Reports** Verify that users can send e-mail with read and delivery receipts, can open and read receipts, and receive non-delivery report (NDR) messages when they send e-mail to users who are not valid.
- **Calendar operations** Verify that users can send and forward both single and recurring appointments and meeting requests, view attendee availability and status, respond to meeting requests with Accept, Tentative, and Decline, view free and busy data, edit series of a recurring meeting request, send meeting updates, and delete a meeting request to remove it from their calendars.

- **Access permissions** Verify that users can send e-mail on behalf of another user (including meeting requests), and that a delegate user can accept a meeting request for another user.
- **Search** Verify that users can search for items in folders, text in messages, messages from or to specific users, messages containing attachments, and so on.
- **Sort** Verify messages sort properly in ascending or descending order and by sender or subject.

Public folders must be part of the client test plan, In addition to the upgrade issues mentioned earlier, even in a native-mode Exchange 2000 organization, be aware of setting permissions on public folders. When an administrator or a client changes the permissions on a public folder. they can create a distribution list and set permissions to allow members of that list to access the public folder. As soon as an access control list (ACL) is used to set permissions on a public folder, all associated distribution groups convert to universal security groups. Not all distribution lists, also called universal distribution groups, convert to universal security groups. This only happens when an ACL is used to set permissions on a public folder. When testing, ensure that all universal security groups are hosted in a native-mode Windows 2000 domain.

Testing considerations for public folders include:

- **Upgrades** Upgrade an Exchange 5.5 server that has a public folder store.
- **Access permissions** Use a public folder with permissions to allow a distribution list to access the public folder. Replicate this public folder to a native-mode Windows 2000 domain. Verify that the universal distribution group converts to a universal security group after it is accessed for the first time. Test public folder replication among Exchange 5.5 and Exchange 2000 servers. Verify content and permissions.
- **Users** Add users from mixed-mode and native-mode domains to the universal security group and verify that Outlook users have appropriate access to the public folders.
- **Outlook** Test various edits to public folder permissions with Outlook.

Third-Party Connectors Test Plan

Incorporate third-party connector and software tests to determine the upgrade strategy to Exchange 2000.

Contact the connector or software vendor to determine connector or software compatibility with Exchange 2000. With this information, you can determine the migration path, if any.

You may not be able to upgrade connectors such as fax, paging, and virus scanners that use the Exchange directory. You may be able to upgrade monitoring and backup software that does not interact with the Exchange directory.

Test upgrades in the test lab individually. If a problem occurs, you can trace it to a specific connector or software.

Test Strategy

Prior to testing, you should establish a strategy that defines specific testing goals and their results and benefits. Create a document that lists features and functions required for the new messaging environment. This list is the basis for new features of Exchange 2000 to test. Determining a strategy with the company's existing messaging environment in mind is important. Also, defining the order of migration and implementation should be one of the goals you work towards.

In addition to testing the new features of Exchange 2000 for general operability, your test lab can have a small-scale model of the existing messaging environment for interoperability and migration testing with Exchange 2000. It may seem obvious that it is important to understand how to install and use a particular feature. However, each tester should keep in mind the level of complexity expected of a given feature in the production environment. While it may be impossible to duplicate the complexity a particular feature may introduce when fully implemented, having the test team fully understand the feature minimizes the risk of difficulties during the actual deployment.

Testing Your Scenarios

Test plans arise from the identification of all of the features to be tested and migration and implementation scenarios planned during the production implementation. A plan that includes taking the first Exchange 2000 server, and its base functionality, into the existing test environment is a good place to start. This may involve a complete build of a new Windows 2000–based server with Active Directory, or an in-place upgrade scenario. Either way, the test plan should involve deployment scenarios expected during implementation.

The list of required features that was a guideline for the test strategy section becomes more important when developing a test phase. End-user and administrator functionality can be grouped into test plans that model how features go into place during implementation.

Exchange 2000 represents a significant change in the functionality and management of the Exchange 5.5 environment. The directory has been moved from Exchange to Active Directory. A user's authentication and Exchange configuration are managed in a single place: the Active Directory Users and Computer snap-in. Exchange 5.5 sites no longer exist, and the Exchange System Manager snap-in replaces the Exchange Administrator program.

Feature Tests

Examine the new features included in Exchange 2000. Each can be tested independently to familiarize the Exchange and Windows 2000 architecture teams and administrators with this new functionality.

Some new features and corresponding test scenarios appear in the following sections. Features such as clustering services and real-time collaboration are not included.

Multiple Databases

Description Allows database size management, larger mailbox size, and faster backups and restores.

Test objective To understand functionality of multiple databases on a single Exchange 2000 server.

Test method Create two mailbox stores, create a user on one store, and dismount the other store. Move user between stores.

Storage Groups

Description Storage groups allow you to assemble databases into easily managed groups. All databases within a storage group share the same transaction log files and log settings. Storage groups and multiple databases enable users to service virtual organizations on one server.

Test objective To understand storage group functionality.

Test method Create a new storage group, create a new mailbox store in the new storage group, and move users between storage groups.

Backup and Recovery

Description Backing up the Exchange 2000 server is a simple process. Each storage group has a set of transaction log files for all the databases in the storage group. You can:

- Back up select databases in a storage group.
- Back up multiple storage groups simultaneously.
- Restore select databases in a backup set.
- Restore databases simultaneously.

A restored single database allows all transaction log files for the storage group to replay, but only the databases being restored will be processed.

Test objective To recover a single database or storage group.

Test method Dismount the database to be restored, restore all incremental and differential backup sets, and then restore the full backup set.

- The administrator selects which databases in a storage group to simultaneously backup.
- The administrator backs up more than one storage group at a time.
- Only the damaged database must be offline to complete a restore. The other database in the backup set or storage group remain online.
- Multiple restore operations can occur simultaneously.

Public Folders

Description Exchange 2000 allows you to create multiple public folder trees.

Test objective To create a new public folder hierarchy.

Test method Create a new public folder hierarchy in System Manager, and then access the public folder hierarchy using Outlook Web Access.

Full-Text Indexing

Description Full-text indexing allows users to search Exchange databases for words or phrases more quickly. Each word in a database is indexed and attachments can also be searched.

Test objective Test functionality of full-text indexing.

Test method Enable and start full-text indexing on both a mailbox and a public folder store, establish an indexing and rebuilding schedule, and search mailbox and public folder stores using an Outlook client.

Address Lists

Description Address lists replaces Address Book Views in Exchange 5.5. The lists are created using a Lightweight Directory Access Protocol (LDAP) query of Active Directory mail-enabled objects.

Test objective Create a new address book list.

Test method In System Manager, create a new address list, and browse the list with an Outlook client.

Routing Groups

Description From a physical, network-layer perspective, routing groups replace Exchange 5.5 sites. Move servers between two routing groups, if both routing groups appear within the same administrative group.

Test objective To understand routing group functionality.

Test method 1 Install Exchange 2000 server in an Exchange 5.5 site, run Active Directory Connector, create a mailbox on the Exchange 2000 server, and then send a message between Exchange 2000 server and Exchange 5.5 server. Install another Exchange 2000 server, create a mailbox on the new Exchange 2000 server, and then send a message to all users on all systems. Verify that the message was received.

Test method 2 Create two routing groups in an administrative group, install an Exchange 2000 server into one routing group, and move the server (drag and drop) into the second routing group.

Message Transport Between Routing Groups

Description There are three ways to send messages between routing groups:

- Routing Group connector
- SMTP connector
- X.400 connector

Test objective To test the functionality of the Routing Group connector.

Test method Create routing group one with servers A and B, create routing group two with servers C and D, create a new mailbox-enabled user and store the mailbox on server B. create another new mailbox-enabled user and store the mailbox on server D, create connections between the routing groups using Routing Group connectors, and send a message between routing groups. Verify receipt of the message. Unplug the network cable from the source bridgehead server A. Resend the message. Verify receipt of the message.

Administrative Groups

Description Different administrative groups are configured to manage Exchange 2000 servers. In mixed mode, only one administrative group is allowed per site. However, there is always at least one routing group per administrative group. The administrative group to routing group ratio is 1:n.

Test objective Test the functionality of administrative groups.

Test method Add a new user by using Active Directory Users and Computers. Assign the new user the Exchange Administrator role using the Exchange Administrative Delegation wizard, log on as the new user, and dismount a store.

Policies

Description Policies allow you to establish limits and rules applicable to databases. Various databases can have different storage limits and retention policies.

Test objective To establish a deleted items retention and mailbox size limitation.

Test method In System Manager, create a new mailbox policy, configure the policy to retain deleted mail for 60 days. Maximize the mailbox limit to 100 MB. Apply the policy to the appropriate mailbox store.

Distribution List Management and Usage

Description Distribution lists are stored in Active Directory. There are two types of groups that can be mail–enabled: security and distribution.

If membership needs to be viewed in all domains in a Windows 2000 mixed-mode environment, create a universal distribution group.

If the group needs to be used for permissions on public folders and includes users from more than one domain, then it must be a universal security group.

Universal security groups must be hosted in a native-mode Windows 2000 domain.

Test objective To create a mail-enabled distribution group with a membership that is visible to all domain users.

Test method Create a new group in the Active Directory Users and Computers snap-in, set group scope to **Global,** set group type to **Distribution,** and then add a user to the new group. To verify, open the address book in an Outlook client; the group should appear.

Web Folders

Description All folders and items within the store are URL-accessible.

Test objective To test URL accessibility to mailboxes and public folders.

Test method 1 (mailbox) In a Web browser, type http:\Lserver\exchange\alins, where server is the Exchange server name, and ***alias*** is a user name. When prompted, enter the user name, domain, and password.

Test method 2 (public folder) In a Web browser, type http:\lrerver\public, where sewer is the Exchange server name. When prompted, enter the user name, domain, and password.

Upgrade Testing

It is possible to perform the migration from Exchange 5.5 to Exchange 2000 in several different ways. The method you choose determines the software upgrade method. The following are typical upgrade paths:

- In-place upgrade from Windows NT Server 4.0 to Windows 2000 Server and Exchange 5.5 to Exchange 2000.
- Move users from an Exchange 5.5 server to an Exchange 2000 server.

Active Directory Connector (ADC) must be installed regardless of the plan to migrate to Active Directory. There are five migration paths for user data to Active Directory:

- In-place domain upgrade followed by the installation of ADC.
- Populate the new Active Directory using ADC.
- Use ADC to populate the new Active Directory structure and then clone users by using Active Directory Migration Tool.
- Use Active Directory Migration Tool with SidHistory, and then run ADC.
- Use Active Directory Migration Tool without SidHistory, re-create the ACLs, and then run ADC.

Evaluate and test each upgrade scenario to determine the best migration path for your organization. For more information on upgrade scenarios, see the Exchange 2000 documentation and other chapters of this book.

CHAPTER 8

Piloting Exchange 2000

Michael Aday, Senior Consultant, Microsoft
Will Martin, Senior Consultant, Microsoft

Before deploying Microsoft Exchange 2000 Server in your company, plan a pilot program to further test and refine deployment strategies and configurations. The pilot program is a scaled-down version of the full deployment, using pilot groups that are representative of the users in the organization.

This chapter addresses questions critical to planning pilot programs. The answers to these questions help you to identify necessary tasks and to prepare resources. such as a lab that represents typical client-server configuration. Ensure your pilot-group volunteers have sufficient time in their schedules and are willing to cooperate in the pilot program.

This chapter also discusses how to proceed once the pilot program is in place: such as meeting with the pilot group to set expectations, maintaining a record of all issues and concerns. Use these records when designing solutions to correct potential problems.

After the pilot program is complete, you can use documentation from the pilot program to create specific guidelines for deploying and configuring Exchange 2000 Server.

In This Chapter

The Role of the Pilot

Determining Pilot Objectives

Who Should Participate?

Setting User Expectations

How Long Will a Pilot Take?

Production Pilots

Moving From Lab to Pilot

Documenting Pilot Processes

Applying Lessons to Production

The Role of the Pilot

After you have tested all the functionality intended for your production Exchange 2000 system, various questions and comments arise, generally along the lines of: Why is there a need to conduct a pilot? All systems worked in the configurations. Why can't we just load them onto our production systems and deploy Exchange?

The following is a short list of selling points for conducting a pilot:

- To document production procedures that are necessary, but are not in place in a lab environment
- To determine typical and expected user interaction with the system before company-wide deployment
- To develop or modify training documentation based on pilot user feedback
- To accurately determine production environment deployment parameters (possibly a phased approach or specific features)
- To determine the level of production support needed for the new application

In most test labs, the entire production configuration is not reproducible at one time. First, the lab tests Microsoft Windows 2000 Server for messaging system connectivity and to identify connections to other messaging systems. Then technicians test Microsoft Windows 2000 domains that contain Exchange 2000 servers. Do not expect these test domains to handle a production system load.

> **Note** Although you can duplicate the load level, do not attempt to duplicate actual user interaction. The production domain presents many issues inherent in a production system, stemming from years of use, and encompassing any hardware glitches discovered over its lifespan.

Lab messaging systems running on Exchange 2000 and earlier versions of Exchange are replicas of what runs in production. Lab systems represent a clean environment rather than a true production environment. A pilot environment connects to actual production systems and domains, so it approximates your actual production environment.

Questions arise during pilot planning: How much of our production system should we involve in a pilot? Do we need all routing group connectors in place in our pilot system'? How about connections to our existing messaging system—do we need to duplicate all of them?

The answer is to conduct a pilot test on as much of your production configuration as possible. At a minimum, test one of each configuration slated for production. If you currently have three Microsoft Mail connections, test one in your pilot. If you have two or more routing groups in your planned production architecture, you need to test at least one routing group connector (and therefore need two routing groups) in your pilot environment. If you plan to host multiple storage groups and messaging databases on a single server, include a server with at least two storage groups and databases.

Pilots happen outside of the production Windows 2000 forest. Once you have completed your pilot, move your mailboxes back to Exchange 5.5, *temporarily* removing the Exchange 2000 environment. Then you can build the deployment topology. Note that a key aspect of moving from the pilot environment to the production environment is using a one-way connection agreement for the Active Directory Connector (ADC). This connection agreement enables you to populate the pilot Active Directory, while being able to remove and reconstruct the environment without jeopardizing your live production. It is unwise to assume that a successful lab deployment means the same for production. Rely instead on a pilot with similar production functionality. This can better illustrate a smooth transition for your messaging and collaborative plans.

Determining Pilot Objectives

Exchange 2000 is a multifaceted application that extends a messaging system's functionality in many ways. It follows that Exchange has many features from which to choose. Companies that define criteria for a successful pilot are positioned to correctly identify milestones or actions to take for the pilot process. Getting to know users' day-to-day messaging activities and visualizing the company's needs can help you to create clear objectives and to clear a path towards determining specific features to include in the pilot.

You must make difficult decisions during the pilot. Specifically, the company needs to decide which features are necessary for the future optimal operation of their messaging and collaboration infrastructure. As each feature wins consensus, its implementation must follow. For example, if the company requires public folders, decisions on folder organization, access, newsgroup usage, and folder creation must be made.

Determining what resources are at your disposal (available participants, risk tolerance level, degree of interoperation, allowable operations overhead), and the perceived long-term value of a successful pilot, all are important steps in the piloting process. Ultimately, the decision to conduct a pilot hinges on a simple cost benefit analysis. Does the cost of the pilot outweigh the benefit of the new features that you want to implement?

Choosing Which Features to Pilot

The new features of Exchange 2000 can be categorized into three main areas.

- **Higher scalability, easier management, and lower cost** Exchange 2000 offers improved scalability that enables customers to continuously expand their system functionality for an increasingly large numbers of users and to address architectural complexities.
  - New storage technologies and multiple mailbox stores and public folder stores per server.
  - User account distribution between mailbox stores, based on availability requirements.
  - Protocol front-end virtual servers.
  - Active Directory Connector, which allows you to synchronize earlier versions of Exchange with Active Directory objects.
- **Greater availability** Exchange 2000 offers increased reliability and reduced downtime after hardware failure.
  - Windows Clustering.
  - Exchange 2000 integration with Microsoft Management Console (MMC) through snap-ins presents the opportunity to customize a subset of Exchange features, such as security, in an MMC console—without starting Exchange System Manager.

    Note This is similar to creating a specialized console to manage a customized set of Windows 2000 services.
- **Enhanced workflow services** Exchange 2000 offers new features to support your knowledge-management strategies. The Microsoft Web Storage System has features that allow various ways for users to communicate. In addition to e-mail, the Web Storage System also supports discussion threads, Microsoft Office documents, Web documents, and voice-mail; it can act as a platform for advanced communication technologies. You can create workflow processes with real-time conferencing, unified messaging, instant messaging, and wireless devices. Use the following features listed below to control information management and workflow.
  - Enhanced Web client capabilities, such as Outlook Web Access.
  - Improved application programming interfaces (API) include new versions of Collaboration Data Object (CDO) programming interfaces.
  - Integrated workflow events, asynchronous or synchronous with store access.
  - Enhanced real-time collaboration services, such as Instant Messaging Service, data and video conferencing, and Exchange 2000 Chat Service.

The process of deciding which features to run through a pilot is universal. Evaluate your direction to determine what features and functions are appropriate for achieving its objectives.

Existing Infrastructure Requirements

Exchange 2000 has new server features (without client software or configuration impact), as well as enhanced services, such as real-time collaboration, that require specific client configurations and supporting infrastructure. This varies with your selected features. However, when you determine infrastructure and client configuration, assess the server- and client-side functionality requirements.

Server-Side Features

To test Exchange 2000 server-side enhancements, identify the degree of interoperability that you want between the test architecture and the production architecture. If the test architecture is intended to be separate from the production architecture, components such as Active Directory Connector configuration agreements may need no configuration. In all cases, consider the following server-side features:

- New storage technologies
- Protocol front-end virtual servers
- Active Directory Connector (if upgrading from Exchange *5.5)*
- Exchange 2000 clustering
- Enhanced Outlook Web Access capabilities
- Improved programmability

All of the listed features are relatively transparent to the client, with no requirement other than a MAPI or Internet protocol-capable client existing on a desktop.

Client-Side Features

Some Exchange 2000 features (such as data and video conferencing, which require multicast support) require that specific versions of Windows exist on the users' desktop. In addition, some features may require infrastructure changes (data and video conferencing require changes to the network routers to enable multicast to span the organization) or specific client-side software on each user's workstation (Instant Messaging Service). Other features, such as enhanced Outlook Web Access services may be accessible with earlier clients (Internet Explorer 4.0 or Netscape Navigator 4.0) but may exhibit different behaviors when viewed with later clients, such as Internet Explorer 5.0.

Who Should Participate?

This section outlines some criteria for determining the number of participants in an Exchange 2000 pilot. These parameters are about identifying aspects that can dominate the pilot. Use the following questions as a guideline:

- Which functions does the company plan to test?
- Which users will have access to the new features; can they test the new functions to the extent they will be used in the production environment?
- What infrastructure supports a specific number of users for the intended feature set?

  Note Infrastructure costs for a pilot feature set are typically much greater than the hardware costs. Consider operational and client support requirements, client-side components, and administrative overhead. These parts make up the whole when deciding on the number of test features. Also, conduct cost-benefit analyses to determine if a function warrants deployment.

- List the potential risks to the production infrastructure when creating the pilot infrastructure. Will a feature endanger client business processes—if so, is the risk worth the benefit?

Depending on the overall scope of implementation, you may need a specific number of users for a pilot. Complex pilots may require more users to generate enough activity to test all the scenarios.

Setting User Expectations

Set expectations of potential pilot participants appropriately. In some cases, it may not be possible to meet stringent production service-level agreements within the pilot architecture. This could be because the software is a pre-release version, or that resources are intended for production infrastructures rather than for a pilot. Discuss the following potential scenarios:

- Availability of pilot services, such as user messaging services. It may be that the user moved to a pilot account without retaining a production account.
- The time it takes to resolve performance and availability issues. This is generally determined by the amount of administrative and troubleshooting resources that are available to address the issues.
- Preservation of user data stored on beta or pilot systems. Even if the systems are maintained as production systems, there is no guarantee against data loss.
- Users involved in the pilot may not be able to see all of the other employees in the company messaging system. This ability depends on the type of directory synchronization between the pilot and production infrastructures.

Carefully evaluate these potential risks for pilot users because they will help you identify future issues for production users.

How Long Will a Pilot Take?

The answer to this question is, "It depends on the company," and this is why it is important to get as much useful information from your pilot users as possible. However, do not allow the pilot to prevent you from deploying the new messaging system. The pilot should bring about the new messaging system faster and more efficiently. For a shorter pilot program, engender frequent and clear communication with your users.

Production Pilots

Many pilots will never be removed from the production environment. Once a pilot starts, it is considered a production system, and treated as a transition from the lab to actual production. Your company determines this, so be prepared to let the pilot stand as originally designed.

Moving From Lab to Pilot

When moving to a production pilot system. ensure that the software on the pilot system is configured as on the tested software. Herein find implementation of the same configurations currently on the lab system (in the pilot environment) along with taking any lessons from the lab installation to the pilot environment.

Hardware configurations can change as long as the relative software configuration remains constant. At no time should software be installed on a production or pilot system that is not lab tested. This includes hardware support software, but if the same hardware is not used in the lab and for the pilot production system, this may not be possible. If so, one of the pilot systems should be temporarily moved to the lab, and the software loaded there before running it on the production system.

Before building a production messaging system, even if it is intended for a pilot environment, document, test, and launch back up and restore procedures for that system. This forms the basis of your production disaster-recovery documentation.

Documenting Lab Configurations

As with any system installation, lab environment software needs modification over time. However, unlike production systems, the lab systems require frequent modifications, and care must he taken to ensure all of these appear in the documented configuration. Build pilot systems that are based on the final lab configuration.

If the lab systems have Microsoft Mail connectors, the base configuration for the pilot system must also contain these connectors. However, if the base lab system does not have them, the final configuration will not.

Obviously, mirror the lab system as closely as possible to the pilot system, but do note any shortcomings of the lab hardware. The lab tests show whether specific hardware platforms are robust enough for the expected pilot and production loads, without actually using those hardware platforms in the test lab.

Part of deciding what should comprise the pilot configuration is based on what was learned in the test lab.

Evaluating Lessons from Lab Tests

Once you have completed lab tests, scrutinize software configuration modifications to determine if the same modification should go into the pilot system. You need to determine if the modifications meet pilot system requirements and the original project plan. Your entire pilot configuration must be based on actual production requirements.

Not all production requirements need to be met by the pilot system. Also, only test functionality in the pilot environment that will be part of the production system.

Documenting Pilot Processes

During your pilot, monitor, administer, back up, restore, and manage your pilot environment. Generate documentation of these processes that will be valuable for the deployment process. You can use it for your final production systems, so add as much detail as possible.

No untested hardware or software configuration should be built on production systems. It is critical that if you plan to use a less powerful system to host your Routing Group connector server in production, you test that system configuration in your pilot before deployment.

Applying Lessons to Production

A pilot determines if a system is ready for production deployment. Consider then, both the pilot documentation and configuration when building the production design. In a well-planned pilot environment, its final system configurations migrate to the first production systems, so all documentation should allow these systems to be duplicated to the last configuration setting.

Take into account that there have been system configuration changes during the pilot. These may be necessary for the system to function properly in your production environment. They therefore must be documented during the pilot to allow you to recreate them in the production rollout. The documentation you create during the pilot becomes the roadmap to a successful production implementation. This means the more you include day-to-day management and configuration in these documents, the easier your job is when you have to reproduce the configurations for your production system.

CHAPTER 9

Preparing a New Environment

Ziad Chbeir, Managing Consultant, Inacom Corporation

Preparing the environment for deploying Microsoft Exchange 2000 Server as a new messaging system requires a plan and a sound architecture. Using a single example, this chapter illustrates how to design and implement Exchange 2000 Server. A fictional company, LitWare, Inc. and its full design for Exchange 2000 will demonstrate the process. For completeness, key design topics will be discussed at a design level and not in depth.

LitWare has 800 retail stores that sell movies and music CDs worldwide. Its headquarters are in New York City and there are shipping offices in Dallas, London, and Hong Kong SAR. The company's worldwide population is 6,000. LitWare has decided to migrate to Exchange 2000 Server from Lotus Notes.

This chapter assumes the following are in place:

- A Windows 2000 environment
- A previous messaging system (Lotus Notes)
- No Exchange system exists
- No Exchange design exists

In This Chapter

Assess and Evaluate Environment

Design Exchange 2000 Architecture

Design Coexistence and Migration Plan

Train Administrators and Users

Design Summary

Assess and Evaluate Environment

Assessing and evaluating the current messaging and network infrastructure consists of the following phases:

- Define the required functionality
- Conduct a gap analysis
- Create the plan
- Prepare the environment

The following sections discuss these phases in detail.

Define the Required Functionality

LitWare's requirements and its expectations for Exchange 2000 Server require evaluation. To assess the required functionality, interview the technical staff and department heads who represent the user population. If the user population is not represented during that phase, you risk designing a system that is not what users want.

Ask questions about business requirements and processes, and how Exchange will contribute to solving existing business problems. This assessment phase is critical to identify user expectations and to intercept potentially expensive mistakes.

LitWare is currently using Lotus Notes for messaging and collaboration. The current messaging infrastructure is used between the main offices only; business needs dictate the need to expand messaging to the retail stores.

Design Goals

The following list of goals is a foundation for designing the Exchange 2000 environment. Exchange 2000 has various design options; these options will be discussed throughout this chapter. The design options that beat match the design goals will be used for the final design. The design goals follow:

- Provide reliable messaging over existing WAN.
- Provide access to all messaging and collaboration functions from all company locations
- Provide e–mail and secure connectivity between all locations.
- Develop applications on the Exchange platform that simplify international shipping and customs issues, and that further integrate workflow concepts into these applications.
- Install real-time collaboration tools that promote effective employee communications, regardless of location.
- Create backup and recovery of Exchange message databases and applications.
- Maintain a single directory system that contains all corporate directory information.
- Implement the system without increasing WAN bandwidth.
- Deliver client performance that meets customer service demands.
- Provide sales access to e-mail from the Internet and offline when needed.

Conduct a Gap Analysis

Conduct a gap analysis with a comprehensive set of questions for LitWare IT employees and decision makers. Ask questions to non-technical department heads who would provide valuable feedback. Your questions should cover the following topics:

- End-user population
- Geographic profile
- Physical network
- Name resolution
- Windows 2000 environment
- Existing messaging environment
- Support

LitWare has 5,000 users in its retail stores, 500 in New York City, 350 in Dallas, 100 in London, and 50 in Hong Kong SAR.

Only one WAN connects the four major offices (New York, Dallas, London, and Hong Kong SAR), and there is no WAN connectivity to the retail stores.

There is a T1 connection between Dallas and New York, a fractional T1 connection between London and New York, a 128–KB frame relay connection between London and Hong Kong SAR, and a 384–Kbps frame relay connection between New York and Hong Kong SAR.

LitWare has already upgraded and migrated to Windows 2000 Server, and is therefore using the Windows 2000 DNS service.

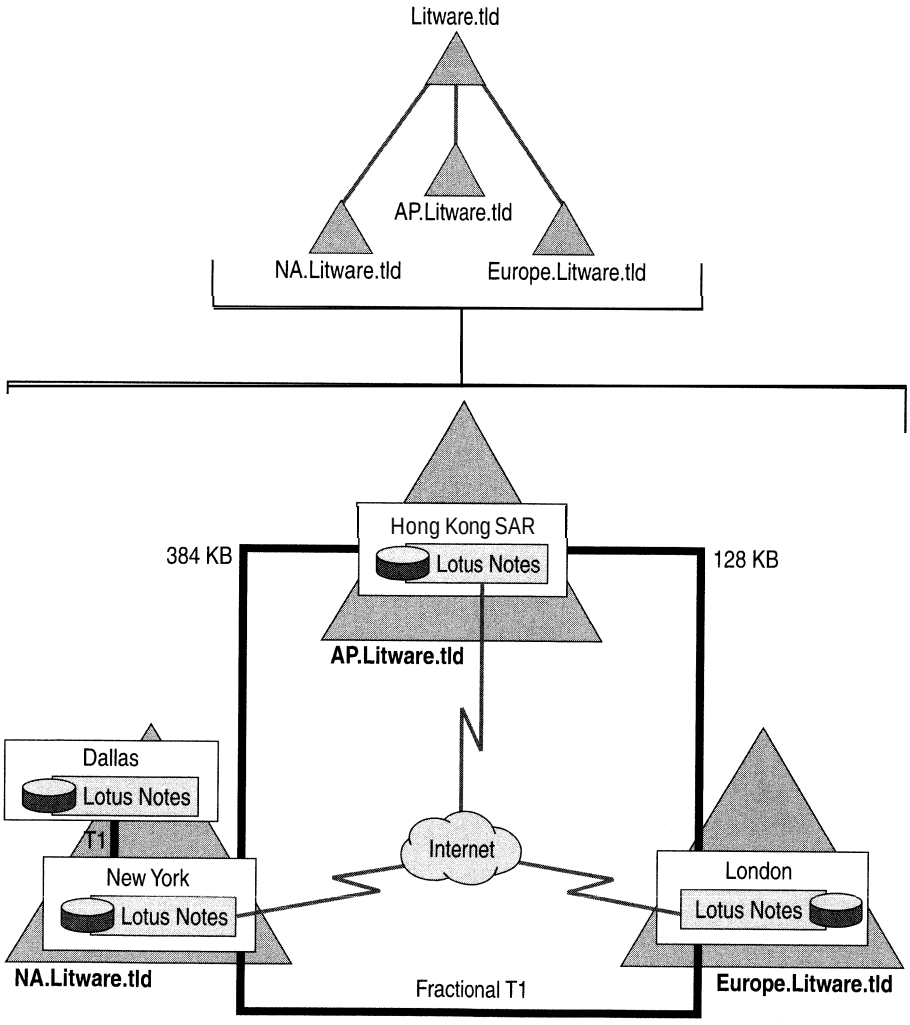


Figure 9.1 Existing messaging topology

LitWare has a single Windows 2000 forest, with four domains:

- us.litware.tld (United States)
- europe.litware.tld
- ap.litware.tld (Asian-Pacific)
- litware.tld placeholder domain

Each of the four locations has a domain controller and a global catalog. New York has two domain controllers.

There are four Notes Named Networks; New York, London, and Hong Kong SAR have gateways to the Internet. Currently there are approximately 1,200 e-mail accounts because the retail stores are not connected. Remote users connect using remote access servers, located in New York.

E-mail usage is moderate, and old data will be migrated from Lotus Notes to Exchange 2000. The current workflow applications will not be migrated. In addition, there are no uniform IT standards. Each location is independently controlled.

Create the Plan

Consider the phases required to produce a design plan for Exchange 2000 that applies the Microsoft Solutions Framework discussed in earlier chapters. This section will not discuss Microsoft Solutions Framework in detail; but it will show how it applies to this fictional installation.

After identifying the system functions and conducting the gap analysis, you need to construct a project plan to provide detailed information on timelines and overall implementation costs.

This approach is based on the four phases of the Microsoft Solutions Framework project methodology: Envisioning, Planning, Developing, and Deploying. This solution is milestone-driven, and focuses primarily on quality assurance and testing.

The project plan will:

- Have a Program Manager, someone with decision-making powers, and who can effectively track the progress of the project.
- Have a defined scope with respect to time, money, and people resources.
- Be up-to-date.

Table 9.1 Microsoft Solutions Framework phases

Phase	Activities	Special Considerations
Envisioning	Strategy, recommendations, high-level risk assessment, concept design, team structure	Consolidate requirements. Define goals and scope based on resources, schedule, and features.
Planning	Analysis, migration and coexistence planning, project planning, migration procedures, standards and policies creation, administrator training	Engender steady communication among those affected by the project.
Developing	Prototype testing, back-end implementation, messaging interoperability solution implementation, pilots, help desk training	Select a pilot group to reflect the users, not just the IT department. The pilot group should accept potential downtimes and design changes.
Deploying	End–user training, rollout, marketing/public relations, ongoing support	Timely availability of proper software, hardware, and technical staff when required is important to meet deadlines and deliver a sound project. Project success is directly dependent on end user satisfaction. Therefore, deployment should focus on minimizing end user downtime and inconveniences.

System functionality should be released in an iterative fashion. For example, it could be decided that a public folders strategy will be implemented in future versions. These system versions can be internally marketed so LitWare's users can know what features to expect and when to expect them. Create a system of versions to control functionality that the IT team can readily and accurately support, while allowing them to accommodate new functionality. Versioning also allows realistic schedules to be met. Deploying too much functionality might overwhelm not only IT, but also the end users.

Table 9.2 Microsoft Solutions Framework project milestones

Project Phase	Milestone
Envisioning	Gather information
	Inventory analysis
	Establish strategy and conceptual design documents
	Identify risks and determine mitigation
Planning	Organize the project team
	Train the project team
	Assess the environment
	Test functionality
	Design the environment
Developing	Test the system
	Train administrators and users
	Conduct regional rollouts

Prepare the Environment

Throughout the design process, and before deployment testing, carefully evaluate the environment to prepare it for the Exchange 2000 implementation. The following sections discuss critical points to consider.

Physical Network

Knowledge of the WAN bandwidth that connects sites is critical in designing an Exchange 2000 implementation. The addition of an application like Exchange that will use the WAN will affect all other network applications. The impact is proportional to available bandwidth and what is required by the new application. Looking at the overall bandwidth, and ignoring the utilization of the link is a common mistake. Monitor network utilization over a period of time to learn the available bandwidth in average, peak, and off-peak times. This data helps with decisions on network upgrade requirements, and when and how to schedule replication.

For example, if the network analysis shows that the network utilization peaks between 9 A.M to 10 A.M., 12:30 P.M. to 1 P.M., and 4 P.M. to 6 P.M., you can avoid scheduling public folder replication at those peak times.

In this illustrated design, LitWare's retail stores do not have WAN connectivity with the main offices. The retail stores use Internet dial-up connections as well as DSL, cable modems, or frame relays. In general, retail stores do not have technical staff; also, the employees do not send e-mail to one another. Most of the e-mail activity is between the main offices and the stores. Moreover, the number of stores might increase or decrease based on market demand, and store locations may also change. Therefore, accessing the e-mail system over the Internet is a practical solution. However, you need to upgrade all dial-up connections to faster, permanent links.

Windows 2000 Global Catalog Considerations

This discussion assumes that the Windows 2000 environment is already designed. Review this design and ensure its compatibility with Exchange 2000. It is possible that the person who customized Windows 2000 did not take Exchange 2000 into consideration. As a result, the system may need fine-tuning to accept the new messaging system.

As you evaluate the Windows 2000 environment, consider the following design points:

- **Forest design** It is important to determine whether all domains are part of the same forest, because Exchange 2000 relies heavily on the Active Directory directory service. An Exchange 2000 organization cannot span multiple Windows 2000 forests. For multiple forests, use multiple Exchange organizations.
- **Domain design** Windows 2000 domain structure is more flexible and scalable than earlier versions; however, it has more influence on Exchange 2000 than Windows NT 4.0 has over Exchange 5.5. In Windows 2000, the domain boundaries define the namespace, and each domain includes one or more domain controllers. When a Windows 2000 domain controller is placed into a Windows NT 4.0 domain, the domain is called a mixed-mode domain. To properly prepare the environment, identify whether the environment is running in mixed mode or in native mode. Do not use universal groups in a mixed-mode environment, because doing so adversely affects the use of security and distribution groups.

Examine the domain tree structure to make sure that Exchange fits into it and that the fit complies with security and administrative requirements. For example, LitWare has four separate domains. Technically, LitWare could have achieved the same results by using a single domain model and organizational units. Litware's domain scheme affects Exchange customization in different ways. One is user management and migration. Moving users between domains requires creating new e-mail-enabled accounts and moving the messages to it. Even though the process is automated, it is not as transparent to users as when users move between servers within a domain. The migration of users from Lotus Notes to Exchange also needs more consideration in a multi-domain environment. For more information about migration, see "Design Coexistence and Migration Plan" later in this chapter.

- **Domain controller and global catalog considerations** A Windows 2000 domain controller holds domain-related data, whereas a global catalog server holds domain-related data and a replica of selected attributes from selected objects. Exchange depends heavily on global catalog servers. While preparing the environment consider the following:
  - Position global catalog servers to enable Exchange to perform quick searches and so users have access to the resources with minimum delays. Have at least one global catalog server per Windows 2000 site. A second global catalog may be required for added redundancy or load balancing.
  - Review which fields are to be replicated to the global catalog servers and update that list based on the design requirement. Exchange users rely heavily on the directory. For example, the Department attribute is not replicated by default. However, there may be a need for the Department attribute to be available for a workflow application that depends on that attribute.
- **Naming conventions** Because Exchange 2000 uses Active Directory, review the Windows 2000 naming conventions. Identify and resolve any naming convention discrepancies. Define and add to the naming convention those attributes that are relevant only to Exchange. E-mail address attributes are an example.

End-User Input

It is crucial to survey end users to understand their business requirements, environment, skills, and training needs. If the system does not meet their needs or the users are not prepared, much confusion and frustration will ensue. Therefore, familiarize yourself with their needs and expectations by comparing the system design to their requirements at each step of the process. After assessing their skills, you can better define the training that they will need.

Finally, a good marketing program is crucial for the success of the project. Inform the users about the coming changes and its benefits before installation starts. Users react negatively to surprises and can become unwilling to let go of the old system.

Design Exchange 2000 Architecture

There are several ways to design Litware's environment. Under each design topic, different possible scenarios (if they exist) will be listed, and compared. The scenario that best fits Litware's design becomes the installation model. Keep in mind that different designs are appropriate for a given environment, and that there is not always a clear right or wrong answer. A correct and complete design meets the overall project goals.

Server Locations

One of the main decisions in the design is planning the location of the servers. LitWare has retail stores that are spread across the world. Careful planning creates simpler and less expensive administration. LitWare has the following options:

- **Option 1** Exchange servers in each of the four main locations and in all of the retail stores.

  All locations have small servers to hold the local users. A global catalog server placed in each location will improve response time when users access the system. In some stores, the global catalog server may run on the same computer as Exchange.

- **Option 2** Exchange servers in the four main offices only.

  The four main offices require large servers to service the retail stores remotely. Each of the four locations needs one or more global catalog servers. Users in the retail stores will access their mailboxes using Microsoft Outlook 2000, Outlook Web Access, an Internet Message Access Protocol version 4 (IMAP4) client, or a Post Office Protocol version 3 (POP3) client, depending on the bandwidth that connects the stores to the main office.

  > **Note** Remember that there are about 5,000 users in 800 retail stores, with an average of six or seven users per store. Most of these users are sales representatives and may have minimal technical backgrounds. The stores do not have WAN connectivity to any of the main offices.

The following table compares the advantages and disadvantages of each option.

Table 9.3 Server locations

Factor	Option 1 Servers in All Locations	Option 2 Servers in the Main Offices
Administration	To manage the servers, a technical person is required for every location, or at least one person per group of retail stores. A technical person for every 10 locations would mean that at a minimum, 80 technical people would be necessary.	IT personnel are required in the four main offices only.
Backup and recovery	Backup and recovery required for all 804 locations.	Backup and recovery limited to the four main offices.
Cost of implementation	Approximately 804 server hardware and site licenses, plus servers for public folders, global catalog servers, connectivity, and so on.	Fewer than 20 servers, however, the servers must be powerful, with multi-processors and multiple RAID storage systems. Clustering is recommended.
Reliability	Not dependent on WAN connectivity. When the link is down, users still have mailbox access. Minimal user impact when a server fails.	No user mailbox access if the WAN goes down. Major user impact if a server fails. However, a cluster solution reduces the chance of system failure.
Response time	Response time is quick. The server is on the LAN.	Response time depends on the WAN utilization and bandwidth.

If you compare the advantages and disadvantages of both designs, adopting servers in the main offices creates a more efficient environment. Use the following table to deploy the Exchange 2000 servers.

Table 9.4 Factors for choosing server location

Number of Users	WAN	Solution
Fewer than 30	64 KB or faster	Remote access.
Fewer than 30	Less than 64 KB, or no WAN connectivity	Remote access. Locations that have no WAN connectivity access their mailboxes over the Internet.
More than 30, but fewer than 300	No WAN	Establish WAN connectivity. Use a small local server* that is also running global catalog services.
More than 30, but fewer than 300	Less than 128 KB	Upgrade the bandwidth (recommended, but not required). Use a small local server* that is also running global catalog services.
More than 30, but fewer than 300	128 KB or faster	Use a small local server* that is also running global catalog services.
More than 300	No WAN	Establish WAN connectivity. Use a large local server*; run a separate global catalog server.
More than 300	Less than 128 KB	Upgrade the bandwidth (recommended, but not required). Use a large local server*; run a separate global catalog server.
More than 300	128 KB or faster.	Use a large local server*; run a separate global catalog server.

*Small and large servers are defined in the "Server Roles" section later in this chapter.

Connectivity from the Retail Stores and the Internet

The majority of the retail stores do not have WAN connectivity to any of the four locations. A flexible solution should allow retail store users to access their mailboxes using one of the following methods: WAN connectivity, an Internet connection, or dialing up to an Internet service provider (ISP).

For retail stores that have WAN connectivity, access to Exchange resources is seamless. Depending en the available bandwidth, users can choose Outlook 2000 or another client. Place user mailboxes on a server that has direct connectivity to the retail stores. Client choices are discussed in the "Client Access" section later in this chapter.

Users who connect over the Internet, either using a permanent or dial-up connection, connect to the front-end servers using Outlook Web Access, an IMAP4 client, or a POP3 client. To permit this access and maintain security, use front-end servers on a perimeter network segment, also called a demilitarized zone (DMZ).

A front-end server is an Exchange 2000 server that does not store user data, but instead, forwards client requests to a back-end server for processing. The front-end server uses Lightweight Directory Access Protocol (LDAP) to query Active Directory to determine on which back-end server the user's mailbox resides. A back-end server is an Exchange 2000 server that maintains at least one information store. This division of function between two servers provides several benefits, particularly in a Web environment. Mainly, front-end servers help reduce the mailbox server load, especially if Secure Sockets Layer (SSL) is being used. Moreover, placing the front-end servers on the perimeter network, and the databases behind the firewall increases protection from the Internet.

Front-end servers will be placed on the perimeter network segment of New York, London, and Hong Kong SAR. Dallas will not have front-end servers because it does not connect directly to the Internet. Initially, each location will have two front-end servers. More servers will be added in the future if the load requires it. To provide adequate security, the global catalog servers will reside on the internal segment and the front-end servers will forward search requests.

If you want POP3, IMAP4, and Hypertext Transfer Protocol (HTTP) or other protocols to work on the front–end server, open ports on both the perimeter network and the internal network firewalls. Open the following ports between the perimeter network and the internal network, and between the Internet and the perimeter network.

- HTTP (80)
- Global catalog (3268) and LDAP (389)
- SSL (443)
- POP3 (110)
- Simple Mail Transfer Protocol (SMTP) (25)
- IMAP4 (143)

The previous list contains the main ports, but does not include all required ports for functionality across the firewall. Obtain the full list of ports during testing.

Tighten LitWare security by configuring the firewall to prohibit communications initiated to the internal network, unless it originates from the Media Access Control (MAC) addresses of the front-end servers. Allow initiation of communication from the internal network to the front-end server so that internal network users can use the front-end servers.

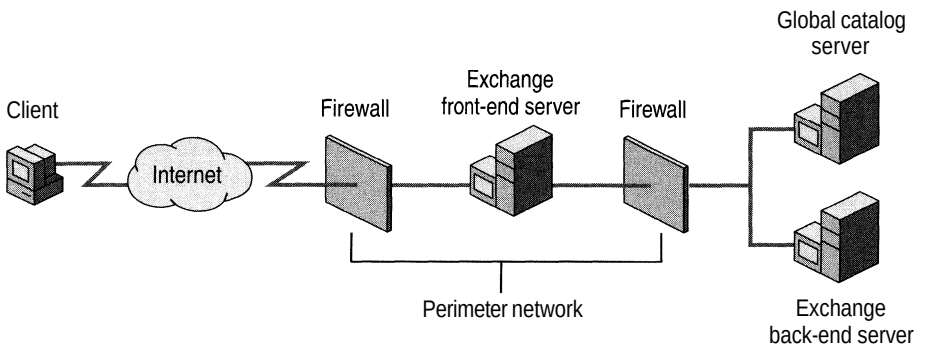


Figure 9.2 Front-end and back-end solution with perimeter network

Global Catalog Placement

To improve performance, place at least one global catalog server in each Windows 2000 site. For added redundancy, two global catalog servers will go in the main offices. For small offices that have fewer than 300 users, the global catalog may reside on the Exchange server to minimize hardware. The New York, London, and Hong Kong SAR offices will each have two independent global catalog servers. The Dallas office gets one separate global catalog server that is not on the Exchange server. If the Dallas office requires more redundancy, add a second global catalog server.

Message Routing

After you have decided on the location of the servers, you need to determine how the servers will connect and how messages will be routed. You can choose from the following options:

- Connect all routing groups by means of a central hub. In the LitWare scenario, the hub goes in the New York City office.
- Create a mesh network, in which each routing group connects directly to all other routing groups.
- Configure routing groups so that their connectivity reflects the underlying network.

Message Routing Option 1

In the first routing configuration, all routing groups connect by means of a central hub, which is located in the New York City office. This configuration is illustrated in Figure 9.3. In this configuration, a routing group placed in New York is a hub and all other routing groups connect to it. Therefore, all messages go through the hub routing group.

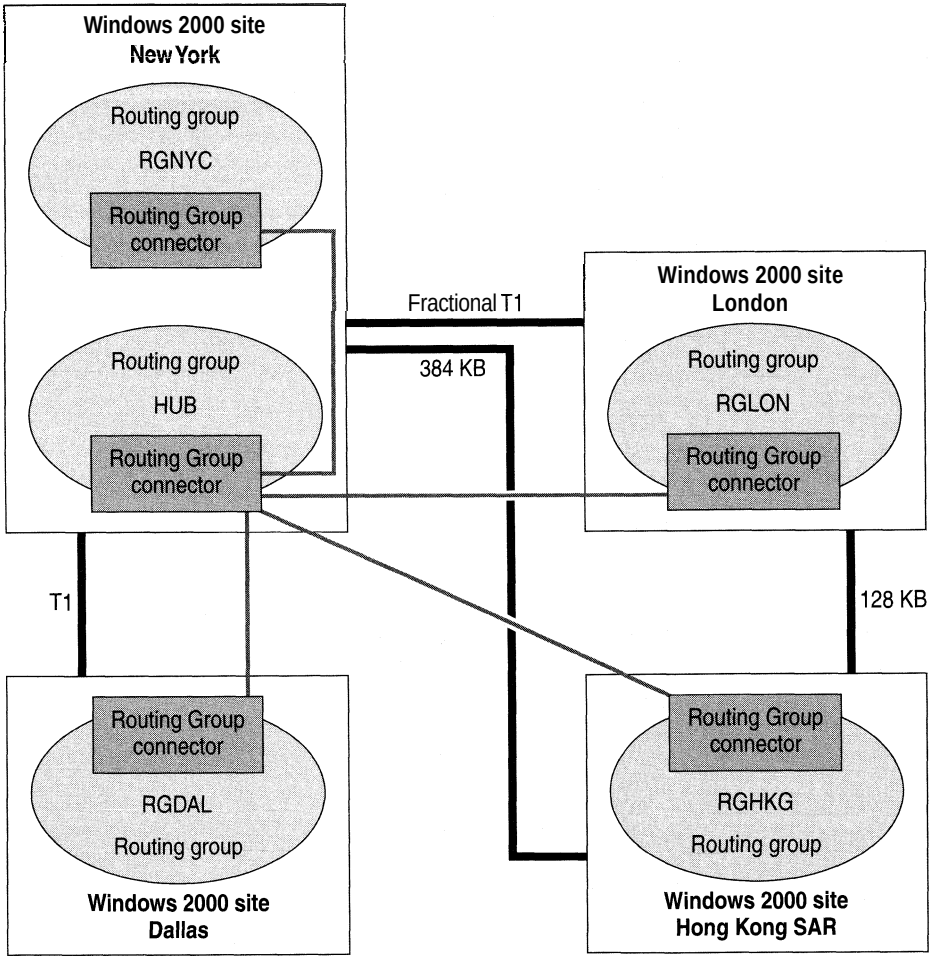


Figure 9.3 Central hub routing group configuration

The central hub routing group configuration has the following advantages:

- All message transfer can be controlled from a single point
- Relatively few connections (four)
- Message tracing is easier

This configuration has the following disadvantages:

- Multiple hops between routing groups
- Message route is not necessarily the shortest WAN path
- No redundant links

Message Routing Option 2

The second routing configuration is a mesh network, in which all routing groups connect directly to each other. In this option, each location has one routing group, and all routing groups connect directly to each other.

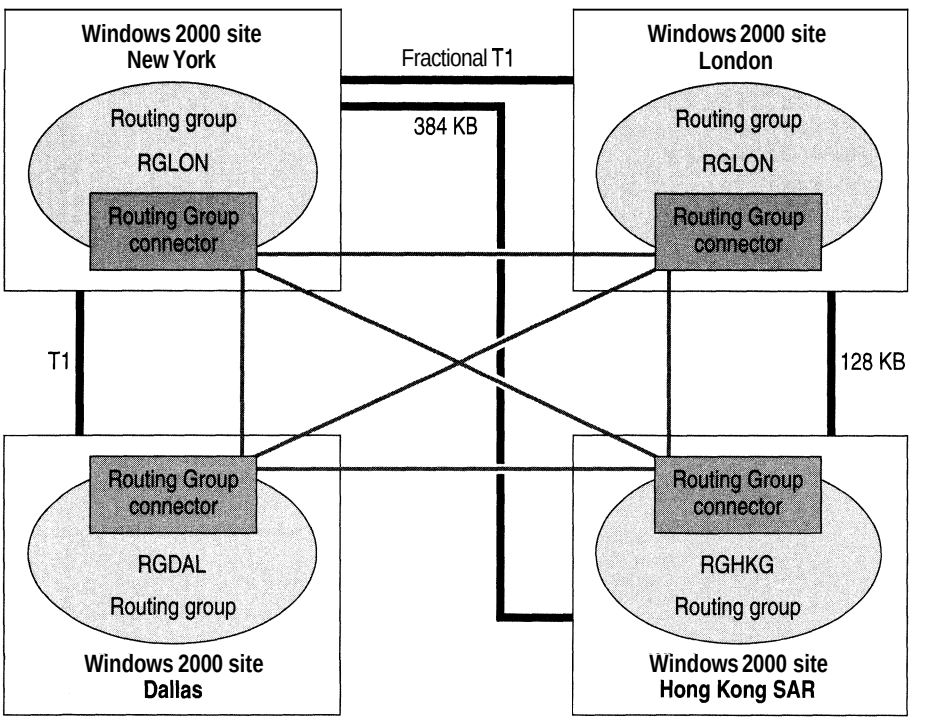


Figure 9.4 Mesh network routing group configuration

The mesh network routing group configuration has the following advantage:

- More connections to monitor and administer

This configuration has the following disadvantages:

- Single-hop routing between routing groups
- Message route does not always reflect the physical path over the WAN

In some cases the redundant links may not be useful when the WAN fails. For example, if the WAN between Dallas and New York fails, all connections from Dallas to New York, London, and Hong Kong SAR become unavailable.

Message Routing Option 3

The third routing configuration reflects the underlying network. Each location is set as a routing group. Connection between the routing groups corresponds to the underlying network topology.

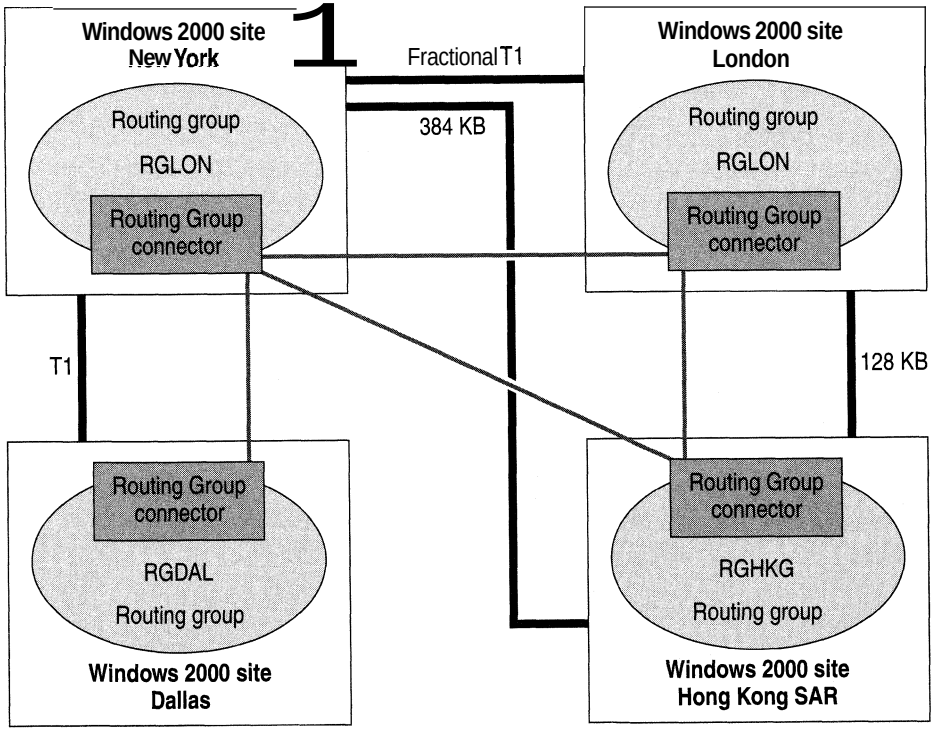


Figure 9.5 Routing group design reflects underlying network

Having the routing group design reflect the underlying network has the following advantages:

- Number of connections directly related to number of WAN links
- Design takes advantage of WAN redundancy

This configuration has the following disadvantage:

- May be harder to trace messages than the "hub" method

If you use the previous comparison, the third option is the best choice because it offers advantages with only one minor disadvantage. Create routing groups according to the following table.

Table 9.5 Routing group design

Routing Group Name	Location	Description
RGNYC	New York	Contains mailboxes and public folders for users that are in the U.S. retail stores and in the New York office. There are about 500 users in New York and 3,500 users in the U.S. retail stores.
RGDAL	Dallas	Contains mailboxes and public folders for users that are in Dallas. This routing group serves about 350 users.
RGLON	London	Contains mailboxes and public folders for users that are in Europe. There are about 100 users in London and 1,000 users in the European retail stores.
RGHKG	Hong Kong SAR	Contains mailboxes and public folders for users that are in Asian Pacific region. There are 50 users in Hong Kong SAR and 500 users in the Asian Pacific retail stores.

In addition to having separate routing groups for each location, place servers on the perimeter network (DMZ) in separate routing groups. Thus, three additional routing groups go in the design: DMZNYC, DMZLON and DMZHKG. These routing groups isolate the perimeter network servers from the internal servers, which restricts and controls communication between the internal servers and the firewall servers.

The justification for these routing groups is as follows: if an SMTP server that resides on the perimeter network is in the *same* routing group as the internal servers (RGNYC for example), the SMTP server sends any inbound message directly to the mailbox server.

Create the additional routing groups on the perimeter network and configure the firewall to allow the SMTP server to access all internal servers in the RGNYC routing group. For tighter security, set up the SMTP server to access only the connector server in RGNYC; allow the connector server to forward the message to the appropriate mailbox server.

Connectivity to the Internet

Three of the four LitWare main offices have independent connections to the Internet. Leave this set up as it is. However, you can consolidate Internet connectivity and establish one point of accesc to the Internet (New York).

Table 9.6 Evaluating Internet access

Evaluation Criteria	Three Internet Connections	One Shared Internet Connection
Distance to connection	One hop	Multiple hops
Number of firewalls	Three	One
Response time	Faster	Slower; all traffic must first reach New York
WAN effects	Minimal; most Internet traffic avoids WAN	WAN links to New York carry Internet traffic
Administration	Three points to administer	Single point to administer

A single Internet connection might improve security because it is more efficient to control one gateway; however, this design affects the WAN, which violates design requirements. Additionally, a single gateway might also affect the response time for the sales staff as they access their e-mail from the Internet. Therefore, Internet connectivity remains distributed when messaging migrates to Exchange 2000.

Real-Time Collaboration Services

Chat Service and Instant Messaging Service do not go in the initial phase of the project. They will be deployed in the future. However, LitWare requests conferencing services during the requirement phase and wants it available to all locations. By using Microsoft Exchange 2000 Conferencing Server, users schedule meetings through the Outlook calendaring interface. The Exchange server starts, manages, and closes the conference.

Data Conferencing Provider is a technology that provides, within an online conference, the facilities to share applications, conduct whiteboard sessions, chat, and transfer documents. There is data collaboration with tools that support T.120, such as Microsoft NetMeeting. Video Conferencing Provider supplies an audio/video facility for each participant. This collaboration is through the Exchange video conferencing client, or NetMeeting.

LitWare intends to make these services available to all employees, but not all locations have adequate bandwidth, so Video Conferencing Provider will not be implemented initially. Only Data Conferencing Provider will be introduced to the users, and voice conferencing (teleconferencing) will occur over a traditional conferencing service. Video Conferencing Provider will be implemented after LitWare upgrades the current infrastructure. This approach satisfies the following design goals:

- **Design Goal 2** Access to all messaging and collaboration functions from all company locations. Employees often move between facilities and stores.
- **Design Goal 5** Deployment of real-time collaboration tools that allow users to work together effectively, regardless of their location.
- **Design Goal 8** Implement the system with no increase of WAN bandwidth.

Data Conferencing Provider

Following are the points to consider when designing the Data Conferencing Provider environment:

- Server locations
- Whether users should be able to join a conference from the Internet, and if so, how to make Internet conferences available and secure
- The number and function of the servers

Conferencing Sewer Location

The servers could be placed in all locations, including the retail stores, or they could be placed only in central locations.

If the servers are placed in all locations, all conferencing features are provided; however, the system is more complex and the cost of network management high. Earlier, it was decided to use centralized servers to simplify management, so the Data Conferencing Provider needs to be centralized. Place Data Conferencing Provider servers in the three main offices: New York, London, and Hong Kong SAR. Dallas will access the servers in the New York office.

Conferencing Internet Connectivity and Security

Decide if Data Conferencing Provider will be accessed from outside LitWare, or only used internally, but ensure that all users have access to data conferencing. Retail store users and traveling salespeople may only have Internet access, so Data Conferencing Provider has to be accessible from the Internet, with the appropriate security measures in place. Because the perimeter network is already secured and has the proper ports open to the Internet, the best place for the Data Conferencing Provider server is on the perimeter network. For users to access the conferencing servers from the Internet, port 1503 for T.120 needs to be open. The servers, by default, do not allow access from the Internet; therefore Internet access needs to be enabled.

Number and Function of Conferencing Servers

Two servers will be deployed in each of the three main offices (New York, London, and Hong Kong SAR). The first server provides both Conference Management Service and Data Conferencing Provider—and the second provides just Data Conferencing Provider. More Data Conferencing Provider servers will be added in the future, if demand on the initial system exceeds capacity. Figure 9.6 illustrates the conference service arrangement in each office. The Conference Management Service server will have a mailbox that contains the conference calendar. That mailbox will reside on the Conference Management Service. System performance will be monitored, and more servers added if necessary.

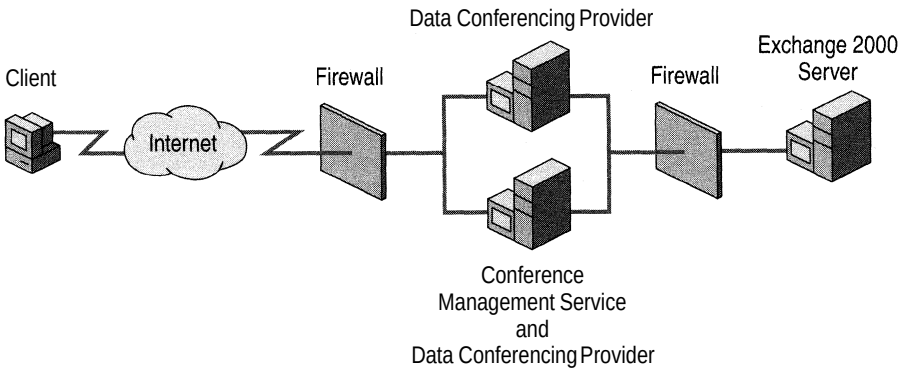


Figure 9.6 Conferencing server deployment

Server Roles

Even though all functionality of Exchange 2000 can be installed on one server, it is recommended to divide the functions across different servers. By doing so, you will accomplish the following:

- **Improved response time** Servers will not be overloaded and the response time will be better.
- **Reduced downtime** Services are not dependent on one another; a failure in one service does not affect other services if they run on separate servers.

LitWare plans to implement the following server roles: mailbox servers, public folders servers, connector servers, front-end servers, and data conferencing servers.

Mailbox Servers

At Litware, each user will have a mailbox limit of 30 megabytes (MB). Each database will have about 36–gigabyte (GB) of storage, using five 9–GB drives with RAID–5 configuration. This allows about 1,000 users per database.

Exchange 2000 can have mailbox stores on the same server, which enables LitWare to have a large number of mailboxes distributed over several databases. This reduces user downtime if a database is being restored after a failure. With earlier versions of Exchange, mailboxes on the same server

reside in the same database. For example, if an Exchange 5.5 server has 2,500 mailboxes, the database may grow up to 75 GB (assuming 30 MB per mailbox), making restoration very time consuming. With Exchange 2000, mailboxes can exist in two separate storage groups, which creates storage groups of about 38 GB each, which will take about half the time to restore.

One option is to have three large mailbox servers. Each server will reside in one of the three main offices (New York, London, and Hong Kong SAR). Each server is for the office and the retail stores that connect directly to it. Thus the New York server has the largest number of users (4,000).

This solution helps reduce the downtime when the server fails, but does not eliminate downtime. To provide high availability, create a cluster solution. With the initial release of Exchange 2000, a two-node cluster is possible. Litware's environment is ideal for a two-node cluster solution. For Litware, one node serves the main office users, and the other node serves the retail stores.

One or more mailboxes, or public folder stores exist within a storage group. Each storage group has one set of transaction log files. LitWare uses two storage groups, one on each node of the server cluster.

For maximum performance, place each database on separate volumes. Moreover, locate each volume on a different controller and data bus. To provide fault tolerance and disaster recovery, place the transaction log files on a volume separate from the corresponding databases. Placing the transaction log files on a separate and dedicated volume increases the performance.

Use the ***C*** volume for the operating system and the Exchange 2000 binary files. This volume should use RAID–1 technology.

> **Note** Because the Dallas server contains only 350 mailboxes, clustering is an option, but not a requirement. Consider clustering for the Dallas server, if the users require constant availability and if e-mail is critical for them.

Table 9.7 Mailbox sewer specifications

Server Component	Large Exchange Server (More Than 1,000 Users)	Small Exchange Server (Fewer Than 1,000 Users)
CPU	Two to four processors	Two processors (upgradeable to four)
Memory	1 GB	512 MB
Operating system and Exchange binary files	9 GB (two 9–GB RAID–1 volumes)	9 GB (two 9–GB RAID–1 volumes)
Network interface	100 MB or to match infrastructure	100 MB or to match infrastructure
Log disk 1	9 GB (two 9 GB RAID–1 volumes)	9 GB (two 9–GB RAID–1 volumes)
Log disk 2	9 GB (two 9 GB RAID–1 volumes)	None
Data storage disks (for each database)	36 GB (four 9-GB and one spare RAID–5 volumes)	36 GB (four 9–GB and one spare RAID–5 volumes)

The following table illustrates the storage system configuration for the New York cluster server. Users in the first storage group will reside on one node of the cluster, and the users in the second storage group will reside on the second node of the cluster.

Table 9.8 Storage system for mailbox servers

Storage Group	Use	Volume	RAID Technology
First storage group	Local mailboxes	E	RAID–5
First storage group	Retail stores mailboxes	F	RAID–5
First storage group	Retail stores mailboxes	G	RAID–5
Second storage group	Retail stores mailboxes	H	RAID–5
Second storage group	Retail stores mailboxes	I	RAID–5
First storage group	Transaction log files	J	RAID–1
Second storage group	Transaction log files	K	RAID–1

Public Folders Servers

It is possible to install public folder servers in a cluster system similar to the mailbox solution. However, public folders are not as critical as mailbox servers, because users are not constantly connected to them. Certain folders may be critical and require continuous availability and that can be accomplished by creating one or more replicas.

Unlike the mailbox servers, the public folder servers use one storage group.

Table 9.9 Storage system for public folder servers

Server Component	Large Exchange Server (More Than 1,000 Users)
CPU	Two to four processors
Memory	1 GB
Operating system and Exchange binary files	9 GB (two 9–GB RAID–1 volumes)
Network interface	100 MB or to match infrastructure
Log disk 1	9 GB (two 9–GB RAID–1 volumes)
Data storage disks (for each database)	36 GB (four 9–GB and one spare RAID–5 volumes)

Connector Servers

Connector servers only route messages; they do not host mailboxes or public folders. A connector server may connect routing groups, run a Lotus Notes connector, or connect to the Internet. A common hardware platform should be implemented for connector servers. This reduces the cost of implementation and future maintenance. Given the size of the company and the relatively low expected load, the servers will not have the optimal configuration for performance. This reduces the cost of hardware. The following table lists a recommended specification standard.

Table 9.10 Connector server specifications

Server Component	Connector Server
CPU	Two to four processors
Memory	256 MB
Operating system and Exchange binary files	9 GB (2 x 9 RAID–1)
Data storage disks	9 GB (2 x 9 RAID–1)
Log disk	9 GB (2 x 9 RAID–1)
Network interface	100 MB

Front-End Servers

Users that access their mailboxes using IMAP, POP3, or HTTP will use front-end servers. This reduces the load of the mailbox servers. Front-end servers will exist in the perimeter network for the sales force and in the internal network for the retail stores.

Table 9.11 Front-end server specifications

Server Component	Front-End Server
CPU	Two to four processors
Memory	1 GB
Operating system and Exchange binary files	9 GB (2 x 9 RAID–1)
Data storage disk(s)	9 GB (2 x 9 RAID–1)
Log disk	9 GB (2 x 9 RAID–1)
Network interface	100 MB

The front-end servers do not host databases, so configure them similarly to the connector server except with more RAM.

Data Conferencing Servers

Use Data Conferencing Provider servers to manage and participate in conferences. These servers do not host mailboxes except for the multipoint control unit (MCU) account; therefore configure them similarly to the connector server. except with more RAM.

Table 9.12 Data Conferencing Provider specifications

Server Component	Data Conferencing Provider Server
CPU	Two to four processors
Memory	1 GB
Operating system and Exchange binary files	9 GB (2 x 9 RAID–1)
Information storage disk(s)	9 GB (2 x 9 RAID–1)
Log disk	9 GB (2 x 9 RAID–1)
Network interface	100 MB

Naming Conventions

When designing an Exchange 5.5 environment, define the namespace before implementation. Namespace is important for Exchange 5.5 because it is static, and a consistent naming convention helps to create an intuitive design. With Exchange 2000, most of the static namespace is defined in Windows 2000, including many user–related fields. This section assumes that most of the namespace has been defined during the Windows 2000 implementation, and discusses the namespace as it relates specifically to Exchange 2000.

For a successful Exchange implementation, it is critical that all naming standards be followed throughout the company.

Exchange Organization Naming

Use the name of the Windows 2000 root domain to name the Exchange organization: LitWare.

Routing Group Naming

In Exchange 2000, the routing group directory name and the routing group display name are the same. You can use either of these naming conventions to set the directory and display names.

- **Abbreviated name** "RG and three characters identifies the geographic coverage (country, region, city, or district). For routing groups that reside on the perimeter network (DMZ), "DMZ and three characters identifies the geographic coverage.
- **Full name** "Routing Group" and the full name of the geographic coverage. For routing groups that reside on the perimeter network (DMZ), "DMZ" and "Routing Group" and the full name of the geographic coverage.

Table 9.13 Routing group name examples

Directory Name	Display Name
DMZNYC	DMZ Routing Group New York City
RGNEWYORKCITY	Routing Group New York City

Administrative Group Naming

Generate administrative group names along the following two guidelines.

- **Abbreviated name** "AG" and three characters identifies the geographic coverage (country, region, city, or district). For example, AGNYC.
- **Full name** "Admin Group" and the full name of the geographic coverage. For example, Admin Group New York City.

Server Naming

Server names uniquely identify each server in the organization. Because sewer names are used by end-users, they should be intuitive. Choose names that indicate the server's geographic location, thereby limiting unnecessary browsing across wide-area links. Exchange 2000 follows the convention for naming servers running Windows 2000:

MSX<function><location><instance #>

<function> Identifies the function of the server, such as MB (mailbox), FF (public folder), BH (bridgehead), and FE (front–end server).

<location> Identifies the location where the server resides, for example, NYC for New York City.

<instance #> These characters provide for multiple servers, per location and function. Single–digit instances require a leading "0".

For example, the first mailbox server in New York City will be named MSXMBNYC01.

SMTP Alias Format

Generate the SMTP alias format for all LitWare employees according to the following guidelines:

- **Standard** Firstname.Lastname@Litware.tld
- **Tiebreakers** Use a middle initial, if necessary, or qualifier (for example, Mike.Smith01@Litware.tld).

This alias format ensures that consistent, simple addressing continues throughout the company and ensures a smooth transition from the existing SMTP environment.

Mail-Enabled Contacts for Non–LitWare Users

Mail-enabled contacts are for outside vendors or customers who access mailboxes from external locations. The display name and the e-mail addresses for mail-enabled contacts appear in the LitWare global catalog.

Display names for mail-enabled contacts for non–LitWare employees follow this convention: *first name, last name (company, role)*. For example, Adam Barr (Airlines International, Travel Coordinator).

Mail-Enabled Contacts for LitWare Users

Use mail-enabled contacts for LitWare users on the earlier e-mail system (Lotus Notes). Keep display names intact after directory synchronization occurs. Names will change according to the accepted conventions during the migration to Exchange.

Lotus Notes Addresses

To coexist with Lotus Notes, define a Lotus Notes proxy address that is compatible with the existing Lotus Notes environment.

The address will be:

FirstNameLastName@Exchange

Defining the Lotus Notes proxy address for all locations eliminates the need to create a new Lotus Notes domain.

The following required information will be exchanged between Lotus Notes and Exchange:

- First Name
- Last Name
- Initial
- Display Name *(last name, first name* in Exchange)
- Company *(division* in Exchange)
- Department
- Title
- Phone
- FAX

Mail-Enabled Groups

Because Exchange 2000 uses Active Directory, the naming standards for mail-enabled groups should follow the Windows 2000 standards.

To reduce administrator over-involvement, assign group ownership to the users that request the group.

Active Directory contains two types of groups: security groups and distribution groups. Members of a security group can be granted access to network resources. Members of a distribution group only exist as a distribution list and cannot receive access to resources because of their membership in a distribution group.

Security Groups

- Grant access to resources.
- Group scopes follow the standards defined by the Windows 2000 design.
- LitWare-wide security groups consist of sub-groups created by the divisions. There are no individual participants.

Distribution Groups

- LitWare-wide distribution groups use universal groups.
- LitWare-wide distribution groups consist of sub-groups created by the divisions. There are no individual participants.
- Distribution lists within a domain use global groups. This creates flexibility if network access needs to be granted in the future.
- Distribution lists that span multiple domains use universal groups.
- Universal groups contain sub-groups that may be universal groups or global groups.

Mandatory Distribution Lists

Create group lists for suppliers and business partners when you set up public folder access lists. Doing this ensures that access is restricted to public folders and that broadcast messages go to these groups, without too much administrative effort.

Create the following:

- An ALL EMPLOYEES universal group
- An ALL NON-EMPLOYEES universal group

Scheduling Resources As Recipients

Generate resource names according to the following guidelines.

- **Display name** –Location–Resource Type
- **Resource types** Conference room (CNF RM), television (TV), LCD panel (LCD).

In the display name, "Location" refers to a physical location, like a building name and floor number.

Using the '~' character (tilde) keeps all resources grouped together, directly after the distribution lists. Location ensures that resources in the same location are listed next to each other.

Virtual Resources

Create virtual resource names along the following guidelines.

- **Display name** –VR Location–Resource Type#
- **Resource types** Virtual team conference (VTC), video conference (VC), teleconference (TC).

Locations are New York, London, and Hong Kong SAR.

Backup and Restore

Develop a backup and restore strategy. The lack of a strategy will severely affect your ability to recover lost data after a disaster.

Use the following guidelines to design a backup and recovery strategy.

- Attempt to standardize company software and hardware; a standard system reduces training requirements for the staff and minimizes the restoration time after disaster occurs.
- Develop and document well-tested backup and recovery procedures. Be sure the appropriate people have access to this documentation at all times.
- Keep a consistent log of the backup tapes and store them in a safe location
- Perform practice recovery on frequent basis and do it on an identical test server. The procedure will help you confirm that your strategy works, your tapes have good data, and you are ready in case of a failure.

Administration

In Exchange 2000, administration is separate from message routing. Divide Exchange administration into logical groups that reflect the administrative structure of LitWare rather than the underlying network. Because the administration of Lotus Notes was not standardized at LitWare, you need to design a new administrative model. The following list contains three administration types to consider:

- **User administration** User account management administrators. Because Exchange 2000 uses Active Directory, integrate user administration with the Windows 2000 user administration team. There will be four user administration groups, one for each main office.
- **Exchange Sewer administration** Exchange Server administrators perform backup and recovery, server upgrades, and so on. There will be four groups, one for each main office.
- **Connection administration** These administrators manage the messaging infrastructure and the connectivity between routing groups, the Internet, and Lotus Notes. This group will have the rights to perform all tasks on any server. However, they will not have the right to manage user accounts.

Public Folders

Execute a public folder strategy before deployment to control use and provide a structured environment where users can easily access information.

Create two initial public folder roots. One highly secured root is for LitWare users. Those users who need to access that root over the Internet are required to use HTTP. This root will be named "All Public Folders." The second root is for public access over HTTP. This hierarchy is the database repository for Litware's Web site. This root will be named "LitWare Public Web Storage."

> **Note** When a single server is set to have both databases, place them on two different RAID–5 volumes. This ensures higher input/output (I/O) performance.

The All Public Folders root will be locked and no users, except the administrator, can create a subfolder or place items in the root. Subfolders are created based on departments, under the root of "All Public Folders." Full control goes to each department head so they can create subfolders under their folder and then delegate permissions.

In addition to the department subfolders, create a "Human Resources" folder. This folder will contain employee information and company policies. All users will have read access to it. Human Resources employees will have full control over this folder.

> **Note** When you restrict the number of users who can create top–level folders, users create a tree that is better organized than when the system is configured to allow anyone to create top-level folders.

The Web administrators can define the structure of "LitWare Public Web Storage." Therefore, secure the root so that only the Web administrator creates objects in it. However, replication strategy and folder affinity will conform to the Exchange public folders design.

In Exchange 5.5, a location identifies a subset of servers in an Exchange site that is connected by a high–bandwidth network. This is needed because sites often include additional remote servers for administrative purposes. In Exchange 2000, defining server location for administrative purposes is not necessary because routing groups define single-hop routing, and administrative groups resolve the administrative issues.

Public Folder Affinity

In earlier versions of Exchange, public folder affinity enables public folder referrals to servers in remote sites and also determines the order in which to refer requests to each of these sites. This has been replaced by enabling referrals between routing groups. The cost for connecting each remote routing group is automatically determined by the cost assigned to the connecter between the routing groups.

Public folder referrals to servers in another routing group are enabled by default when a routing group connector exists between the two routing groups. The routing group connector is unidirectional, requiring two instances to be configured for bi-directional traffic. Public folder referrals can be configured for the routing groups in each direction and are transitive.

Servers hosting the public folders will reside on the internal network. They will be accessed from the Internet, through referral by front-end servers.

Each of the main offices has public folder servers. These servers have the data related to the location and a replica of the "Human Resources" folder. A user who needs to access a public folder that resides in another routing group can do so because Exchange makes referrals on the user's behalf. When a public folder is used frequently between multiple locations, establish a replication so that referral across WAN connections is less frequent. Consider the public folder hierarchy below:

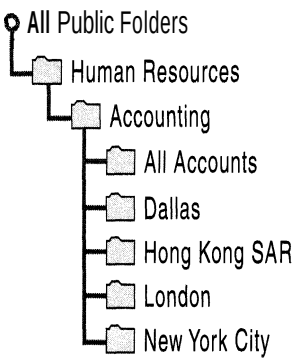


Figure 9.7 Public folder hierarchy

In Figure 9.7, under the Accounting folder, each office hosts its own subfolder. The All Accounts subfolder will be replicated to all four main offices.

The frequency of the replication depends on the type of data and how frequently it changes. In the Human Resources folder, the subfolder holding the employee manual does not change frequently; thus, it is safe to replicate this folder once a day, at night. On the other hand, the subfolder that contains daily news may need to be replicated more frequently because it can be updated many times a day. A frequency of two hours is adequate.

Review domain trusts to ensure adequate client access. Because client computers must use remote procedure call (KPC) to connect to public folder replicas, the client must be authenticated before it can access the data. Study your domain setup to ensure that clients can make connections to any servers that maintain public folder replicas.

Future Public Folder Considerations

Public folders play a very important role in information sharing and collaboration. In a well–organized public folder hierarchy, users have easy access to information. The power of public folders is that they are accessible using a browser. Additionally, public folders can take advantage of Exchange messaging functionality to provide workflow solutions, such as a sales automation solution where traveling salespeople can enter an order on a Web-enabled public folder to automate a customer request. Many other workflow solutions may also be considered, such as automating the expense report process, or developing a solution that manages employee information.

LitWare may consider the use of the Digital Dashboard, a dynamic Web page running within Outlook 2000. The Digital Dashboard consolidates personal, team, corporate, and external information with single-click access to analytical and collaborative tools. It brings an integrated view of a company's knowledge source to a user's desktop, providing immediate access to needed business information.

Digital Dashboards solve the problems associated with information overload by pulling together key information sources into a consolidated view. For more information about Digital Dashboard, see the Microsoft Web site at http://www.microsoft.com.

Client Access

The following section discusses the situations in which different clients can be used. For information about clients, see "Branch Office Scenarios" in this book.

Project design goals must be kept in mind when selecting clients. The following are goals to consider:

- **Design Goal 2** Access to all messaging and collaboration functions from all company locations. Employees often move between facilities and stores.
- **Design Goal 4** Develop applications on the Exchange platform that simplify international shipping and customs issues, and further integrate workflow concepts into these applications.
- **Design Goal 5** Deployment of real-time collaboration tools that allow users to work together effectively regardless of their location.
- **Design Goal 9** Client performance should be at an acceptable level for users to access information on customer demand.
- **Design Goal 10** Sales staff needs to access e-mail from the Internet and be able to use their e-mail offline.

Outlook 2000

After evaluating client features, it is clear that Outlook 2000 meets the design goals for users that have WAN connectivity, and therefore it will be the main client in the four main offices and in any location that has WAN connectivity.

Outlook Web Access

Outlook Web Access provides the user with an Outlook experience while using a browser. Even though Outlook Web Access works on most browsers, use Internet Explorer 5.0 for best performance. Outlook Web Access users can access most of the Exchange services, including their contact information, calendar features, public folders, conferencing services, and more, regardless of the platform that they are using. However, Outlook Web Access does not provide offline functionality. Outlook Web Access meets all of the requirements of the design goals discussed in this chapter, except for design goal 10, which requires offline access. Therefore, Outlook Web Access will be best suited for those users at the retail stores and for the traveling sales people.

POP3 and IMAP4

IMAP4 and POP3 are Internet standard protocols for retrieving e-mail from the mail server. POP3 clients are small and can only access the Inbox. IMAP4 has more features and maintains messages on the mail server while enabling users to maintain duplicates of their folders and messages. Both protocols provide an offline solution for reading and sending messages. However, IMAP4 provides the capability to access and manage more than one folder and access the public folders. When users have the choice, IMAP4 client is preferred over a POP3 client.

The traveling sales people who need offline e-mail access will use IMAP4, and Outlook Web Access whenever they need to access their calendar, contacts, or conferencing services.

NetMeeting

To participate in a data conference, users need a T.120–compliant application such as NetMeeting. Versions of NetMeeting that work with Data Conferencing Provider include NetMeeting 2.0 through NetMeeting 3.0*x*. However, NetMeeting 3.01 and later support secure conferences by encrypting their connection into the associated T.120 conference. For LitWare, NetMeeting 3.01, or later is the choice because most LitWare users will access their conferences from the Internet.

Design Coexistence and Migration Plan

LitWare is using Lotus Notes version 4.6 as their messaging system. To migrate to Exchange 2000, LitWare may consider one of the two options: phased migration or one-step migration.

In a phased migration, the Lotus Notes environment coexists with Exchange 2000, and users are migrated according to location or department. In a one-step migration, all users are migrated at the same time with no coexistence required. Even though a one-step migration saves administrators the effort of planning for coexistence and maintaining two environments, it may be difficult when it is applied to many distributed users. The challenge is that migrating LitWare requires a sufficient number of support staff to help the users during and after migration. Therefore, a phased migration is the preferred solution. The migration strategy divides into two overlapping phases: coexistence and migration. Each phase has a set of recommended tools and approaches.

The current design has four Notes Named Networks, each residing in one of the four main offices. There are three points of access to the Internet: New York, London, and Hong Kong SAR.

As part of the coexistence strategy, the Internet gateway is migrated to Exchange first

Coexistence

There are two reasons for the coexistence phase:

To establish temporary coexistence between Exchange and Notes as migration occurs.

- To stop any further installation of Notes within the company that would hinder or prolong the migration.

Use the following services for the coexistence phase:

- **Mail connectivity** Allows regular e-mail messages to go between Exchange and Notes. Mail connectivity should include support for extended e-mail forms like Telephone Messages, Tasks, and Routing Slips not just those for Message and Reply.
- **Directory synchronization** This bi-directional synchronization of Active Directory and Notes directories allows users complete and accurate address information to identify other users on either system. Directory synchronization is especially critical during migration because addressing information is volatile at that time.
- **Replicate discussion and document distribution databases** These are widely used Notes applications. Their Exchange and Outlook counterparts are public folders. They should be replicated by using public folders during migration.

These services allow Exchange and Notes users to communicate with each other regardless of which system they are using.

Coexistence services deliver the following benefits:

- LitWare users who have migrated to Exchange are not cut off from users, information, or applications they had access to in Notes. They have no incentive to revert to using Notes.

  LitWare users who have not yet migrated to Exchange can communicate with their colleagues who have migrated to Exchange. This transparency may remove doubts or fears about accomplishing work after they migrate to Exchange or Outlook.

Coexistence Architecture

Consider domain architecture, user account issues, and gateway configuration options when planning coexistence and migration.

The Windows 2000 domain provides security boundaries; therefore, keep in mind certain limitations. The limitation that has an impact on the design is the fact that during directory synchronization between Lotus Notes and Exchange 2000, the Lotus Notes recipient only exists in one Windows 2000 domain. Ideally, the Lotus Notes contact information should be placed in the Windows 2000 domain where the users will migrate.

For the users that are in Europe, Lotus Notes contacts need to be placed in the europe.litware.tld domain; the users in the U.S. need to be placed in the na.litware.tld domain, and so on.

If the users do not go in the proper domain, delete the contact information and re-create it in the correct domain during the migration process. This is manageable when working with contacts; however, it may become more complex if the Lotus Notes entries must match Windows 2000 enabled user accounts. Because most of the users in the four main offices already have Lotus Notes and Windows 2000 accounts, it is more practical to map the Lotus Notes account to the Windows 2000 account instead of adding contacts for each Lotus Notes account to Windows 2000, which would double the number of users in Active Directory. Where the gateway is in the na.litware.tld domain, map the Lotus Notes data to the accounts that are in the na.litware.tld domain. The other two domains will not be able to map their user account to the appropriate Lotus Notes entry.

LitWare has two potential solutions for this problem. The first is to keep one gateway placed in one domain, and to repopulate the directory during the migration. The second option is to create three Lotus Notes-to-Exchange gateways, each in one of the three Windows 2000 domains. The Lotus Notes servers in LitWare become three separate environments, based on the three continents and main offices. Each server connects to the gateway nearby. Separating the Lotus Notes environment at LitWare is relatively simple because it already has four Notes Named Networks.

The second solution makes it simple to manage and migrate to Lotus Notes and will work well because Exchange works with backbone messaging. However, consider the following:

- The Lotus Notes directory will be down during the transition. The users in the U.S. will lose the entries of the users in Europe and Asia until they get them back from Exchange. The same will happen for the Europe and Asia environments. Therefore, define expectations for the users.
- Lotus Notes local address books on the client computers will need re-synchronization because the directory will change.
- Explain to decision makers that Notes-to-Notes traffic will go through Exchange. Exchange functions as a robust backbone, therefore lab environment tests can help.

After concerns are resolved, separate the Lotus Notes environment to provide a smoother coexistence and migration.

Three Lotus Notes to Exchange gateways will be set, each in one of the three routing groups: RGNYC, RGLON, and RGHKG.

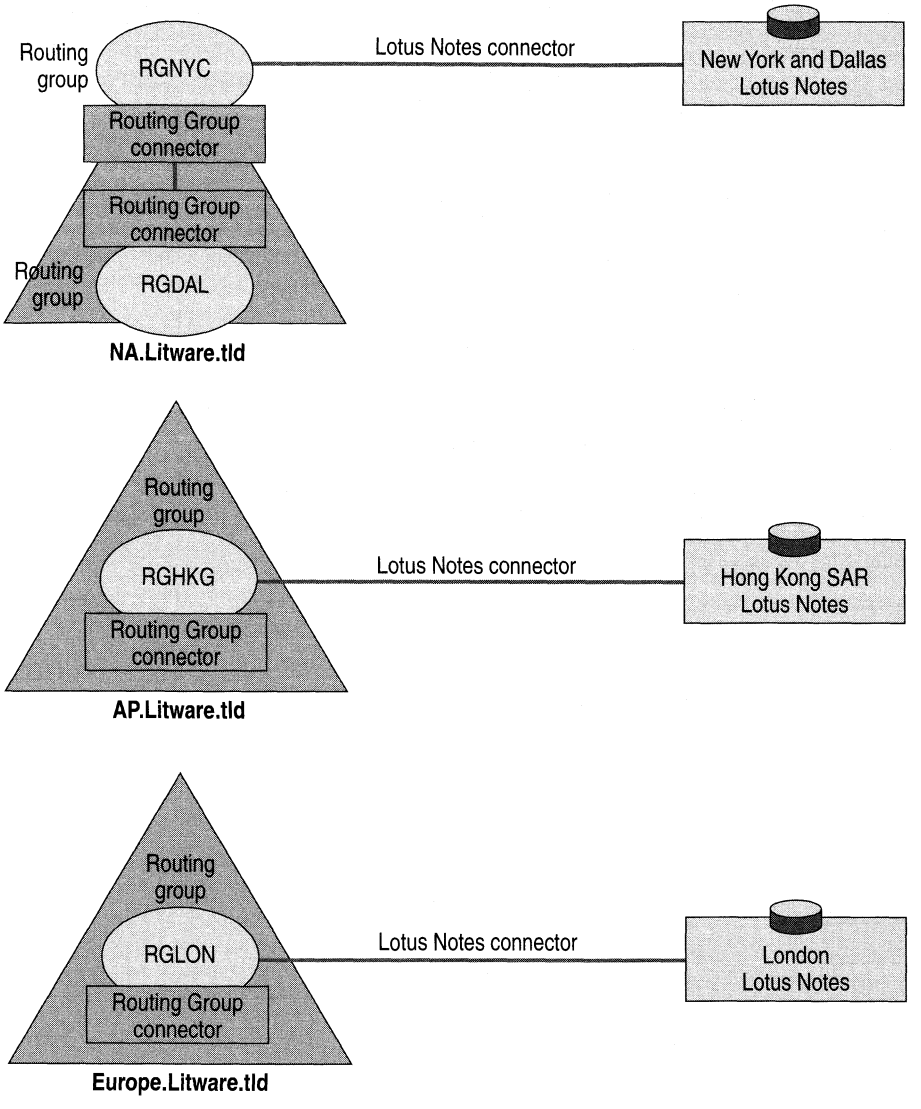


Figure 9.8 Coexistence topology

After you have decided where to place the gateways, decide if the synchronized Lotus Notes entries will be stored in Active Directory as contacts, user accounts, or synchronized with existing accounts. When setting up directory synchronization, the following options are available:

- A disabled Windows 2000 user account
- A new Windows 2000 user account
- A Windows 2000 contact

Most of the Lotus Notes users are expected to have a Windows 2000 account already, therefore the directory synchronization may be set to create Windows 2000 contacts. The Windows 2000 contact will then be merged to the appropriate user account, by using the Active Directory Account Cleanup Wizard.

When setting up the connection between Lotus Notes and Exchange, configure the gateways. The following are aspects to consider:

- **Notes Doc-link conversions** You can convert doc-links to OLE documents, Rich Text Format (RTF) attachments, or URL shortcuts. If all the Lotus Notes servers in the environment are Domino servers and are configured to support Web-enabled applications, convert Doc-links to URL shortcuts; otherwise accept the default, which is RTF attachments.
- **Fields mapping** Using mapping rules, the connector enables you to map Lotus Notes fields to Active Directory attributes and vice versa. You should identify beforehand, which fields you would like to synchronize.
- **SMTP addresses** Generate the appropriate SMTP addresses on the Lotus Notes contacts in Active Directory. This is important because the SMTP gateway will be moved to Exchange and Lotus Notes recipients will expect to continue to receive e-mail that is sent to their current SMTP addresses. Create an appropriate mapping rule during setup.

Migration

Migration, from a messaging perspective, is the process of moving the Notes messaging system to Exchange. Litware's migration involves making copies of existing mailboxes. messages, calendar, and other data, and Importing that information into Exchange. Migration also consists of moving any personal archives, personal address book information, and migrating the Lotus Notes distribution lists to Exchange. Design the migration to Exchange to keep downtime to a minimum. The tools for the LitWare messaging migration are:

- **Migration Wizard** This tool has two components: the source extractor and the migration file importer. The source extractor copies directory information, messages, calendar information, and collaboration data from Lotus Notes and saves the information into a file. The migration file importer imports the created file into Exchange. You can run both of these tools in one step, or separate each function.
- **Active Directory Account Cleanup Wizard** Because the Migration Wizard may create duplicate accounts in Active Directory, this tool simplifies removing duplicate accounts. This tool may not be needed during migration if the coexistence plan was successfully executed and Active Directory was properly prepared in the previous phase.
- **Personal Message Archive Importers** The Outlook **2000** Import feature and the Lotus Notes Mail Importer may be needed.

LitWare's approach to migration is multi-phase. This allows LitWare to move users in different groups separately while testing and optimizing the migration process. This approach also reduces downtime.

Special Considerations

During the migration process, Lotus Notes databases move to Exchange public folders for which you use the Application Converter for Lotus Notes tool. However, when migrating complex applications that may require extensive reprogramming, consider keeping the application on the Lotus Notes server. This is a good option if the application is Web-enabled. By keeping the application on Lotus Notes, users maintain one client. However, in the case of LitWare, keeping an application on Lotus Notes requires some additional planning to allow user access from the Internet. Fortunately, LitWare has no Lotus Notes applications to migrate.

Coexistence and Migration Process

The coexistence and migration process will follow these steps:

- Connecting the gateways
- Configuring the Internet backbone
- Migrating

Connecting the Gateways

The first step is connecting Lotus Notes to Exchange, and ensuring that the e-mail flows smoothly between the systems. A gateway will go in North America, Europe. and in Asia. LitWare must choose whether to connect the three gateways simultaneously, or do the connections in three steps.

The benefit of connecting the gateways simultaneously is that users are only interrupted once. However, contingency planning may be difficult. If users are required to revert to their initial environments, LitWare would need to drop the Lotus Notes directory that was created through Exchange, and rebuild it after connecting the three Lotus Notes environments.

The three-step process may take longer, but it has advantages. During the first step, users are not affected and notice no changes. This is because the first step only connects Notes to Exchange over one gateway, and Exchange has an empty directory. After that connection is made, LitWare has some time before the next step. The next two steps require that Notes Named Networks be disconnected from the current Lotus Notes environment and reconnected using Exchange.

Configuring the Internet Backbone

After the Lotus Notes environment connects to Exchange, LitWare can move the SMTP gateway from Lotus Notes to Exchange. The purpose of moving the Internet gateway before the migration is to guarantee a stable environment for migrated users, and to avoid scheduling a downtime after migration. In addition, the Exchange SMTP connector is a more reliable gateway than the Lotus Notes SMTP message transfer agent (MTA). LitWare improves Internet e-mail for the Lotus Notes users by moving the gateway to Exchange in the early phases.

Migrating

After the backbone is stable, migration occurs. After users migrate, the environment will be cleaned up. The Lotus Notes servers will be decommissioned and the connectors removed from the Exchange servers.

Train Administrators and Users

Incorporate a training schedule into your plan. Training will be required for administrators and the end users.

Administrator Training

A training program should be designed to train the administrators and to discuss the migration process and timelines. After the administrators have been educated and become comfortable with the system, they will transfer their knowledge to the help desk staff, who performs the actual installation.

After the administrators have been trained and have a familiarity with Windows 2000, discuss the following topics:

- Overview of Exchange 2000
- Group and user account management (add, remove, change)
- New server set up
- Server maintenance
- Backup and recovery
- Maintaining public folders, resources, and conferences
- Connectivity to Notes and between groups and the Internet

These topics may be separated into two separate training sessions. The first to help administrators who deal with user accounts; the second to educate administrators who manage the server infrastructure.

End-User Training

During migration, users could be given written material that introduces them to Outlook. This newsletter will explain the basics of e-mail, as well as how to schedule resources, meetings, and conferencing services. The newsletter will also list policies and where to get more information and help, if needed. Moreover, the deployment staff can assist the users, and answer any questions. Following are the methods that may help the users to become familiar with Outlook:

- **Newsletters** Keep users informed. provide basic steps to common procedures, and contact information.
- **Training classes** Conduct training classes on the different clients. Have classes cover all clients that are available at LitWare.
- **Books** Have reference material available for users.

Design Summary

The following figure summarizes the design decisions made throughout the chapter and shows how the systems interconnect. Note that only the New York office appears in the illustration. Mention to users that any design is dynamic and changes in technology and the environment are inevitable. Therefore, a successful design provides for future changes and grows with the evolving environment.

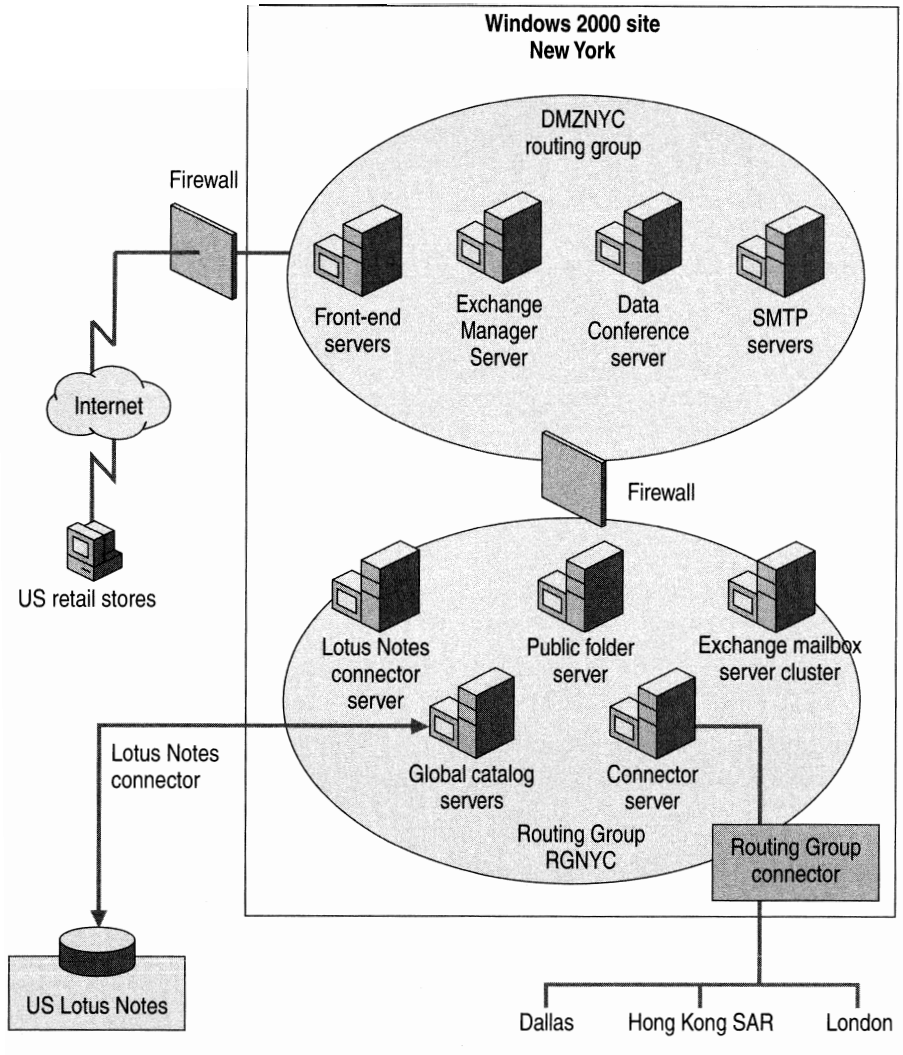


Figure 9.9 Design summary

LitWare may someday connect all retail stores with a high bandwidth WAN connection that will reduce their dependency on the Internet. Figure 9.9 shows separate servers for the Lotus Notes connectors. That is the recommended practice for isolating messaging functions. Because these connectors are temporary. LitWare may configure the Routing Group connector to provide Lotus Notes connectivity or keep a separate server for Lotus Notes. After Lotus Notes is decommissioned, the Lotus Notes connector servers may be transformed into Routing Group connector cervers for added redundancy.

CHAPTER 10

Preparing an Existing Environment

Andrew Holmes, Consultant, Microsoft

As with any complex interconnecting system, the options, considerations, and limitations associated with upgrading an existing Exchange solution can seem daunting. However, the key to a successful deployment is your preparation. This chapter discusses what to do to prepare your existing Exchange environment for the deployment of Microsoft Exchange 2000 Server.

During the upgrade process, there will be several times when decisions you make cannot be reversed. Therefore, it is important early in the process to gather up-to-date and accurate information about your system. This enables you to make informed decisions about your upgrade and migration.

If your existing environment is complex, one upgrade strategy may not work for every part of the system, so you may need different strategies for different servers or Exchange sites. The information you collect about your system and the organization helps determine which strategy to use. In "Installation Considerations" later in this chapter you will look at each component of Exchange and consider the effects of upgrade and migration.

In This Chapter

Operating System Dependencies

Exchange 2000 Requirements

Installation Considerations

Operating System Dependencies

The first step to working out the best upgrade and migration route is to bring existing system information up to date. Auditing the environment is an important part of this.

In many ways, Exchange 2000 has a greater dependency on the operating system than earlier versions of Exchange. This is because many components of Microsoft Exchange 5.5 are now part of Microsoft Windows 2000. Individual dependencies are discussed later in this chapter. To fully understand how this affects your company, meet with your Information Technology (IT) group to communicate the requirements Microsoft Exchange 2000 places on Windows 2000 architecture. These fall into five key areas:

- Exchange 2000 uses Windows 2000 Active Directory directory service to store its schema, naming, and configuration data.
- Exchange 2000 administration is managed through the Windows 2000 Microsoft Management Console (MMC).
- Windows 2000 Internet Information Services (IIS) now provides the lnternet transport protocols for Exchange 2000.
- Windows 2000 Domain Name System (DNS) service provides the location system that allows computers to automatically discover each other on a network.
- Exchange 2000 no longer uses sites. However, Windows 2000 uses sites to logically structure computers on high capacity networks, and Exchange 2000 uses these Windows 2000 sites for message routing decisions. There is no direct correlation between Exchange sites and Windows 2000 sites.

Exchange 2000 Requirements

Exchange 2000 has a number of infrastructure requirements and dependencies to meet when planning an upgrade to an existing Exchange environment.

Plan carefully for the hardware you plan to introduce or re-use. Not only must the configuration of CPU, memory, and disk be sufficient for the role each server is to perform in Exchange 2000, but these servers must also be on the Hardware Compatibility List and be supported by Windows 2000.

Prior to installation of Exchange 2000, your Windows 2000 domain must run DNS and Active Directory. The Active Directory schema must be extended using ForestPrep, and each domain prepared using DomainPrep. Each coniputer that will be running Exchange 2000 needs Network News Transfer Protocol (NNTP) and Simple Mail Transfer Protocol (SMTP) services installed. Before you upgrade. your existing Exchange server must be running Microsoft Exchaiige Server version 5.5 with Service Pack 3 (SP3) on Windows 2000 Server.

Hardware Requirements

The recommended minimum hardware (which exceeds the minimum required hardware) for an Exchange 2000 server is a single Pentium II 300 MHz processor with 256 MB of RAM and 4 GB of hard disk space. However, hardware technology is constantly improving; what you could purchase yesterday can be improved upon tomorrow for the same cost. Base your choices about configuration of hardware on the role each server will play and the load it must support. Your lab testing will help determine what to deploy.

For example, you currently have an Exchange 5.5 mailbox server deployed for optimum disk performance. The private information store would be placed on a separate redundant array of independent disks (RAID) 5 drive array, and transaction log files would be stored on a RAID 1 (or RAID-1+3). In Exchange 2000, deploying a similar mailbox server but with multiple mailbox stores, optiinum performance would be achieved by placing each mailbox store on a separate RAID 5 or RAID 0+1 drive array. But. every set of transaction log files would need to be stored on a separate RAID 1 (or RAID l+3) per storage group.

If your existing server hardware is not a high performance server, you can deploy that hardware elsewhere for another purpose and upgrade to new hardware by using the leapfrog upgrade method. (With the leapfrog upgrade method, you introduce a new server with Exchange 2000 while your existing servers continue to support your e-mail users. When the new server is functioning properly, it becomes the primary mail server.) Alternatively, you can move mailboxes from a computer running any version of Exchange to another computer that is sufficiently configured.

You will also need to ensure your hardware is present on the Microsoft Hardware Compatibility List, which is available on the Microsoft Web site at http://www.microsoft.com.

Windows 2000 Server Infrastructure

Because Exchange 2000 is closely integrated with Active Directory and many former Exchange services are now part of the Windows 2000 Server operating system, there are a number of special requirements for preparing the operating system for Exchange 2000.

Compatible Versions

Exchange 2000 installs only on Windows 2000 Server, Microsoft Windows 2000 Advanced Server, or Windows 2000 Datacenter Server. If you are planning to re-use existing hardware currently running an earlier version of Exchange on Microsoft Windows NT Server 3.51 or on Microsoft Windows NT Server 4.0, upgrade to Windows 2000 Server before you upgrade Exchange.

Exchange 2000 requires four components of Windows 2000 Server:

- Active Directory
- DNS service
- NNTP service
- SMTP service

These must be installed before upgrading to or deploying Exchange 2000.

Active Directory

Exchange 2000 relies on Active Directory to provide directory services. Earlier versions of Exchange use a directory service specific to Exchange.

Each Windows 2000 domain must contain a domain controller. In a Windows 2000 forest where Exchange 2000 servers are installed in more than one domain, at least one server in each domain must be configured as a global catalog server.

The number and placement of domain controllers and global catalogs depends on your network and messaging system topology. For more information about deploying domain controllers and global catalogs, see "Active Directory Design" in this book.

Active Directory Schema

When you first install Windows 2000, Active Directory does not have the structure in place to store information that is specific to Exchange. This structure is called the Active Directory schema and you cannot install Exchange 2000 without extending it. The following sections discuss how you can prepare the schema.

The architecture behind Windows 2000 security allows delegated rights to be given to administrators and users at different levels. To prepare Windows 2000 for Exchange 2000 and to install Exchange 2000 you must belong to the Schema Admins and Enterprise Admins security groups and have administrator privileges on the local computer. Many organizations do not have messaging administrators with all of these permissions. Therefore, use the ForestPrep and DomainPrep tools to separate the tasks requiring Schema Admins and Enterprise Admins permissions from the installation tasks that require local administrator permissions. After the preparation tasks are complete, messaging administrators can install Exchange.

You don't need to run ForestPrep if all the following conditions are met:

- All Exchange 2000 servers are in a single domain.
- That single domain contains the schema master.
- All Exchange users are in that domain.
- The account installing Exchange 2000 belongs to the Windows 2000 Enterprise Admins and Schema Admins security groups.

ForestPrep and DomainPrep run as command line switches of Setup.exe—with the commands `Setup.exe /ForestPrep` and `Setup.exe /DomainPrep`. ForestPrep is run once per forest to extend the Active Directory schema, to nominate the Exchange 2000 administrators, to create the Exchange 2000 organization object in the configuration–namingcontext, and to set up the permissions structures. DomainPrep is run once per domain to create the container for the public folder proxy object and to set permissions within the domain. You do not need to run DomainPrep in a domain until you are ready to install Exchange 2000. Run DomainPrep in all domains where you will install Exchange 2000 and in all domains that contain recipient objects, such as users, mailboxes, and distribution lists.

ForestPrep

To run ForestPrep, you must:

- Log on with Windows 2000 Enterprise Admins and Schema Admins permissions and local computer Administrator privileges.
- Have network connectivity to the server that is the schema master.
- If you are joining an existing Exchange *5.5* organization, run from an account that also has at least View-Only Administrator permissions on the Exchange *5.5* site container and configuration container.

  Important When you update the schema with ForestPrep, you cannot reverse the changes you make to Active Directory.

To join an Exchange 2000 server to an existing organization, install Exchange *5.5* SP3 and the Exchange version of Active Directory Connector (ADC) before you run ForestPrep. The version of ADC that comes with Windows 2000 is not sufficient. Supply the server name of an Exchange *5.5* SP3 server within that organization and the password of the Exchange *5.5* service account.

ForestPrep asks for the account with Exchange Full Administrator permissions. The person or group responsible for installing Exchange uses this account to perform subsequent Exchange 2000 installations.

After the first installation of Exchange 2000, this user or group can create Exchange administrator accounts throughout the forest by using Exchange Administration Delegation Wizard. The Exchange Full Administrator account is different from the Exchange *5.5* service account the wizard requests if you are installing into an existing site. This account is used by earlier versions of Exchange *5.5* to start Exchange services and for server-to-server authentication. ForestPrep uses this service account to access information about the organization in the Exchange *5.5* directory.

There are other considerations you must take into account when installing Exchange 2000 into an existing Exchange 5.5 organization. If you are logged on to the schema master domain and you specify an existing Exchange 5.5 server with SP3 in a Windows 2000 child domain during ForestPrep, you have Full Exchange Administrator rights in the child domain as well.

ForestPrep uses the information you provide to extend the Active Directory schema to include information that is specific to Exchange. This information is located in the files Schema0.ldf–Schema9.ldf on the Microsoft Exchange 2000 Server installation CD in \Setup\i386\Exchange. Around 700 classes and attributes will be modified or added. Exact details are available in the Exchange 2000 Software Development Kit (SDK) at the Microsoft Developer Network (MSDN) Web site at http://msdn.microsoft.com/library.

The new objects and changed attributes are replicated forest–wide to all domain controllers in the schema-naming context. Exchange attributes are also added to the partial replica set, which is a partial replica of the information within each domain that is replicated to global catalog servers. This partial replica contains a subset of the attributes of all directory objects and is read–only. Because Exchange attributes are added to the partial replica set, a complete replication of the global catalog occurs when you have more than one domain. Depending on the size of Active Directory, it can take a considerable amount of time to replicate throughout the forest. If your forest will grow considerably in number of servers and objects in the directory, run ForestPrep early in the deployment process. Subsequently, only changes are replicated between global catalogs. This is more efficient than the replication method used by Exchange 5.5. In Exchange 5.5, a complete record replicates even if only one attribute has changed. With Exchange 2000, replication only occurs only between Windows 2000 domain controllers instead of between all Exchange 5.5 servers.

ForestPrep grants necessary permissions for new Exchange 2000 administrator accounts. This prevents Domain Admins and Enterprise Admins from obtaining Exchange 2000 privileges through inheritance of extended privileges further up the directory service hierarchy. It is therefore important to give the Exchange 2000 Administrator permissions to users and groups by using the Exchange Administration Delegation Wizard. Exchange organization objects are added to Active Directory, and rights are assigned to the Exchange Administrator account for those objects.

DomainPrep

DomainPrep is a tool that you can use to prepare the forest without performing an Exchange installation. After running ForestPrep, run DomainPrep in each domain where you want to install Exchange 2000, including the domain in which you ran ForestPrep. Also run DomainPrep in each domain that contains recipient objects, such as users, mailboxes, and distribution lists. To run DomainPrep, your account should be a member of the Domain Adrnins security group for the domain and have local administrative permissions on the computer from which you are running DomainPrep. Exchange 2000 and Exchange 5.5 administrative permissions are not required to run DomainPrep. After running ForestPrep, you must wait for replication to complete in a domain before running DomainPrep in that domain.

The procedure for running DomainPrep is the same whether or not your forest includes earlier versions of Exchange.

When you use DomainPrep, it will do the following:

- Create the global security group, Exchange Domain Servers (formerly known as Domain EXServers).
- Create the domain local security group, Exchange Enterprise Servers (formerly known as All Exchange Servers).
- Add the Exchange Domain Servers group to the Exchange Enierprisc Servers group.
- Create the container for the public folder proxy object.
- Grant appropriate permissions for Exchange 2000 administrators and Exchange servers.

The first Exchange 2000 server installed in the domain becomes the Recipient Update Service by default.

DNS Service

At least one server in the domain must provide the Windows 2000 DNS service, configured to allow dynamic updates.

NNTP Service

The NNTP service, provided with Windows 2000, must be present on each server before you install Exchange 2000.

SMTP Service

In Exchange 2000, the SMTP service, a component of Windows 2000, provides enhanced functionality. SMTP is now the default transport protocol for Exchange and must be present on each server before you install Exchange 2000.

Optional Windows 2000 Services

If you plan to allow dial-up access, install Windows 2000 Routing and Remote Access. This service is used by Exchange 2000 Server to connect to the Internet by using a modem, or by clients requesting remote access.

Exchange 4.0, Exchange 5 and Exchange 5 pendencies

Upgrading or migrating resources to Exchange 2000 places requirements on the existing Exchange deployment. This section discusses what must be in place to upgrade to and coexist with Exchange 2000.

Upgrade Paths

Your current Exchange environment may include various versions of Exchange with different service packs. Upgrading a server to Exchange 2000 requires that your server run Exchange 5.5 SP3. Exchange 2000 Server includes an update for Exchange Server version 4.0 and Exchange Server version 5.0 that have specific service packs installed. For more information about upgrading to Exchange 5.5 SP3, see the Exchange Web site at http://www.microsoft.com/exchange.

Migration Options

If existing hardware does not meet Exchange 2000 minimum requirements, you can use the Exchange 5.5 Move Server Wizard to move resources between servers in an Exchange site. Sites containing Exchange 5.5 and earlier versions should have the latest service packs installed.

Installation Considerations

Take the following considerations into account as you prepare to upgrade your existing Exchange 5.5 organization.

Disk Space

An in-place upgrade, which refers to an upgrade of a server rather than an upgrade of an Exchange organization, requires an additional 30 percent of disk space on the partition that contains public and private information stores.

Time Required to Upgrade a Server

The process of upgrading the private information store and public information store to Microsoft Web Storage System is complex. Tests performed at customer sites and in Microsoft labs have shown that fast disks and processors affect upgrade times. Your own test lab enables you to gauge the amount of time your servers need to complete an upgrade. Testing can also help you prepare for potential issues with your upgrade process and to plan for contingencies.

An upgrade to Exchange 2000 is faster than earlier upgrades because data upgrades are deferred until a user requests the data for the first time. Deferred data upgrades reduce the time for setup, and extend the conversion over a longer period of time. Data retrieval time is slightly longer for the first access to data, but the time difference is usually imperceptible to the user.

Permissions Required to Upgrade

After preparing your organization with ForestPrep and DomainPrep, the account you specified as the Exchange 2000 administrator has Exchange Full Administrator permissions. When this account has been added to the Administrators group for the computer, the account can install the first instance of Exchange 2000 without having to be a member of the Enterprise Admins, Schema Admins, or Domain Admins groups. By using the Exchange Administration Delegation Wizard, this account can also delegate to other users or groups various Exchange administrator roles, including the ability to install Exchange.

Windows Trusts

If you plan to install new Exchange 2000 servers into an existing Exchange site, and your current Exchange 5.5 sewers are in a Microsoft Windows NT Server 4.0 domain, you must configure the Windows NT Server domain to trust Windows 2000. This enables you to add the new Exchange 2000 administrator account with service account administrator rights throughout the Exchange 5.5 directory. Otherwise, errors will occur during setup.

Server Roles

The order in which you upgrade your servers to Exchange 2000 is not important. You can upgrade the infrastructure servers first (public folder and connector servers, for example) and mailbox servers second. If you need to upgrade mailbox server hardware, you can upgrade mailbox servers first. Exchange 2000 can support either method of deployment.

Mixed and Native Modes

The installation of the first Exchange 2000 server sets the msExchMixedMode attribute on the Exchange organization object in Active Directory. If the first Exchange 2000 server joins an existing Exchange site, this attribute is set to true and Exchange 2000 runs in mixed mode. This is consistent with the constraint of a mixed–mode organization, as the routing group membership is the same as the parent administrative group.

When the first Exchange 2000 server installed does not join an existing Exchange site, the msExchMixedMode attribute is still set to true. To enable full flexibility in routing groups and administrative groups, switch from mixed mode to native mode. When you switch to native mode, you cannot include any Exchange 5.5 or earlier versions of Exchange in your organization, nor can you switch back to mixed mode. To switch to native mode, in System Manager, right–click the organization, click **Properties,** click the **General** tab, and then click **Change Mode.**

When an Exchange 5.5 server exists in the Exchange 2000 organization, the organization is in mixed mode. Exchange 5.5 does not have the ability to separate routing groups and administrative groups, and the Exchange 5.5 site functions as both groups. To operate in a mixed–mode environment, each Exchange 2000 administrative group should have at least one routing group that contains Exchange 2000 servers.

Routing and Routing Masters

Each Exchange 2000 server has a map that is represented in a link state table and that describes the messaging topology of the routing group. This map is regularly updated and propagated to all the servers in the topology so that each server can determine not only the most efficient route for a message, but whether the connectors that make up the route are functioning. Link state information is most effective when multiple routing groups are configured in an organization, particularly if redundant paths are available. Each routing group has a master, and the master maintains link state information from different sources and propagates that information to the rest of the servers. In Exchange 2000, the master is, by default, the first server installed in a routing group, though you can change this through System Manager. The network operates less efficiently when the master is not working, therefore attempt to make the master available to all servers in the routing group at all times.

Every Exchange 5.5 site has a routing master, which by default is the first Exchange 5.5 server in the site. This routing master maintains the gateway address routing table. When an Exchange 2000 server joins the site, the administrator must decide whether the routing master should be an Exchange 5.5 server or an Exchange 2000 server. The only time the routing master should be an Exchange 2000 server is if an existing Exchange 5.5 routing server is upgraded or the administrator manually sets it through the Exchange 5.5 Administrator program. The gateway address routing table always replicates to Active Directory in a mixed–mode administrative group as an object called **LegacyGWART**. Routing masters monitor the **LegacyGWART** object and one exists in every administrative group. Until the topology migrates closer to the Exchange 2000 routing scheme, an Exchange 5.5 server ought to remain the routing master. When there are more Exchange 5.5 servers and connectors than Exchange 2000 servers and connectors. you need to leave an Exchange 5.5 server as the routing master.

For more information about how to determine which Exchange 5.5 (or earlier version) server is the routing master, see Knowledge Base article Q235396, "How to Determine the First Exchange Server Computer in the Site" on the Product Support Services Web site at http://support.microsoft.com/directory/. For more information about how to change the routing master server, see Knowledge Base article Q152959, "How to Remove the First Exchange Server in a Site" on the Product Support Services Web site.

Some existing Exchange 5.*x* organization\ use sub-sites, or *Locations,* an Exchange 5.*x* concept that Exchange 2000 does not recognize. Therefore, the gateway address routing table generated by Exchange 2000 would be incorrect. To mirror functionality of sub-site routing, it is possible in Exchange 2000 to add more routing groups to an administrative group in a mixed–mode organization. These additional routing groups can only contain Exchange servers that are in the same administrative group.

Offline Address Books

Exchange 2000 Setup does not upgrade an Exchange 5.5 offline address book server if there are other Exchange 5.5 servers in the administrative group. This prevents an offline address book folder on an Exchange 2000 server from existing and never being updated, which could confuse users who download offline address books from the Exchange 2000 server. If you want to upgrade your Exchange 5.5 offline address book server while other Exchange 5.5 servers are in an administrative group, move the Exchange 5.5 offline address book to another Exchange 5.5 server.

Offline address books are organization–specific in Exchange 2000, not site–specific as in Exchange 5.5. There is no synchronization of offline address books between Exchange 5.5 and Exchange 2000. Exchange 2000 generates its offline address book from Active Directory. Mailboxes on Exchange 5.5 servers use an offline address book generated by an Exchange 5.5 server. The Exchange 2000 default offline address list has an **Exchange 4.0 and 5.0 compatibility** checkbox. When you select this checkbox, Exchange 2000 supports Microsoft Exchange Client 4.0 or Microsoft Outlook (the version that released with Exchange Server 5.0) connecting to an offline address book on an Exchange 2000 server.

Schedule+ Free/Busy and Organizational Forms Library Public Folders

The Schedule+ Free/Busy Public Folder and the Organizational Forms Library Public Folder should replicate in native–mode and mixed–mode Exchange 2000 environments. These folders replicate the same way in Exchange 2000 as they do in Exchange 5.5.

If you do not replicate these folders to other servers, there might be problems when you remove the system folder server. In this case, you need to recreate the system folders on another server. This requirement is the same in Exchange 2000 and Exchange 5.5.

Exchange Component Comparison

This section examines the components of Exchange 5.5 and analyzes what you should consider when you upgrade servers to Exchange 2000 or install an Exchange 2000 server into an existing site and migrate resources to it.

Table 10.1 shows which Exchange 5.5 features upgrade to Exchange 2000 features.

Table 10.1 Upgraded Exchange 5.5 features

Exchange 5.5 Feature	Upgraded	Comments
Lotus cc:Mail connector	Yes	
IBM Office Vision connector	No	
Lotus Notes connector	Yes	
SNADS connector	No	

Table 10.1 Upgraded Exchange 5.5 features *(continued)*

Exchange 5.5 Feature	Upgraded	Comments
MS Mail connector	Yes	
Schedule+ Free/Busy connector	Yes	
MS Mail Directory Synchronization	Yes	
Address book views	No	Replaced by address lists and build rules for address lists in Exchange 2000.
Dial-up connections	No	
Custom recipients	Yes	Requires ADC
Distribution lists	Yes	Requires ADC
Recipients container	Yes	Requires ADC
Mailboxes	Yes	Requires ADC
IMAP4	Partial	Only values in the mailbox and public folder stores, not values in the configuration
POP3	Partial	Only values in the mailbox and public folder stores, not values in the configuration
TCP/IP MTA stack	Yes	
RAS MTA stack	No	
TP4 stack	No	
X.25 Eicon Stack	No	
TCP/IP X.400 connector	Yes	
RAS X.400 connector	No	
Directory structure	Yes	Directory structure after upgrade remains the same.
Public folders	Yes	
System folders	Yes	
Chat Service	Yes	
Key Management Service	Yes	
Outlook Web Access	Yes	
Event scripts	Yes	Might require additional work

Directory Objects

There is a major shift in the conceptual definition of the user and the task of user management between Microsoft Exchange 2000 Server and earlier versions of Exchange. This is because Exchange 2000 mailbox management is integrated with Windows 2000 Active Directory user management. Consequently, certain key terms need to be redefined.

Mailbox-Enabled and Mail-Enabled Objects

Exchange 2000 makes a distinction between a mail-enabled object and a mailbox-enabled object. A mail-enabled object is an Active Directory object with at least one defined e-mail address. A contact is a mail-enabled object because it has a defined e-mail address.

A mailbox-enabled object is an Active Directory object that has one or more Exchange mailboxes associated with it. In Exchange 2000, only a user object can have a mailbox. Thus, to be mailbox-enabled, an object has to be a Windows 2000 *security principal* (that is, the object can log on).

A user is a Windows 2000 security principal that may have Exchange e-mail addresses, or an Exchange mailbox, or both. Thus, a user can be either mail-enabled or mailbox-enabled.

Custom Recipients

Custom recipients are recipient addresses or mail location addresses outside the Exchange organization. In Exchange 2000 there is no longer an Exchange directory, only Active Directory. Use the Active Directory Connector (ADC) to move custom recipients to Active Directory where they become contacts. When you create a contact in Active Directory, you do not need to immediately specify an Exchange 2000 e-mail address. Instead, you can add an e-mail address to a contact after you create it.

Distribution Lists

Like custom recipients, distribution lists migrate to Active Directory. The migration of distribution lists depends on your Windows 2000 domain topology. For more information about distribution lists, see "Distribution Lists and Security Groups" in this chapter.

Recipient Containers

Like custom recipients, recipient containers no longer exist in Exchange 2000. This functionality exists in organizational unit containers within Active Directory.

Address Book View Containers

Microsoft Exchange Server version *5.0* introduced the concept of address book views. Address book views enable the administrator to expose different views based on field groupings to Outlook users. With Exchange Server 5.*x*, Outlook users see each site, its containers, and the aggregated global address list when they look in the server–based address lists. For example, to create a virtual container of all users in a sales team, the administrator creates a view grouped by the **Department** field. Although virtual containers work well, they have certain limitations. For example, if you create a view based on a field, a virtual container is created for every unique instance of data within that field. This means that you would automatically create team containers in addition to the single container you might have intended to create. This issue exists because the rules that determine how the view is created are not flexible enough for larger organizations.

Address Book Views

Address book views become Exchange *2000* address lists. The degree to which you can specify and create lists has improved, and the method for specifying how lists appear has changed. Therefore, if your current address book views are not upgraded, re-create them using the new Lightweight Directory Access Protocol (LDAP) filter rules.

Hidden Objects in the Exchange 5.5 Directory

If you have Exchange *5.5* and Exchange *2000* installed in the same administrative group, and you have an ADC connection agreement configured, hidden objects that are replicated in Exchange *5.5* are visible in System Manager.

Using the Exchange *5.5* Administrator program, objects can be marked with a hidden attribute. To view these through the Exchange *5.5* Administrator program, on the **View** menu, click **Hidden Recipients.** Exchange *2000* allows you to hide objects from the user view also, but there is no special view in System Manager. Objects that replicate from Exchange Server *5.5*, marked as hidden, remain hidden from the user view, but are visible in the Exchange *2000* System Manager.

> **Important** Do not administer Exchange *5.5* hidden distribution lists from Active Directory Users and Computers. Instead, manage them from the Exchange *5.5* Administrator program.

Names of Exchange 5.5 Custom Attributes

When you view Exchange 2000 objects through Active Directory Users and Computers, view custom attributes by clicking the **Custom Attributes** button on the **Exchange Advanced** tab. Although you can populate custom attributes, the names of the attributes do not change. This is by design. Custom attributes in Exchange *2000* are included for backwards compatibility with Exchange Server *5.5*. Therefore, if replication with Exchange Server *5.5* is taking place, the default names of the custom attributes may be changed. To create modifiable custom attributes, extend the Windows *2000* schema by using the Active Directory Service Interfaces (ADSI) editor (Adsiedit.msc). To install this tool, run **Setup** from the Windows *2000* Server or Windows *2000* Advanced Server installation CD in the \Support\Tools directory.

Web Storage System

With the introduction of Microsoft Web Storage System in Exchange 2000, considerations for planning data slorage and replication changes must he made. This section identifies some of these changes.

Mailboxes

Mailboxes in Exchange 5.5 become mail–enabled Active Directory users in Exchange 2000. The users must exist in Active Directory before an in-place upgrade occurs, or Exchange places the mailbox object in the mailbox store's deleted items location.

Move Mailbox Function

The Exchange 2000 **Move Mailbox** function in Active Directory Users and Computers fails if the LDAP port on an Exchange 5.5 server is not port 389. Avoid this by moving mailboxes to another Exchange 5.5 server in the site with a 389 LDAP port, and then use the **Move Mailbox** function to move users.

Public Folders

Permissions and administration for public folders are more complex since the introduction of multiple database support in Web Storage System. The following sections cover new considerations for Exchange 2000.

Public Folder Permissions

A major change in public folders in Exchange 2000 is with access control lists (ACLs). In Exchange 2000 all ACLs are based on Windows 2000 user accounts, instead of Exchange mailboxes. During an upgrade from Exchange 5.5 to Exchange 2000 and during subsequent interoperation with Exchange 5.5, Exchange 2000 servers attempt to match ACLs on mailboxes to ACLs on user accounts, and vice versa. However, both of the following are possible:

- ACLs on Exchange 5.5 mailboxes do not match any Windows 2000 users.
- ACLs on Windows 2000 user accounts do not match Exchange 5.5 mailboxes.

The first possibility is more common because the old Exchange 5.5 ACLs can have users that have been removed from the directory. To avoid this, run the Directory Service/Information Store consistency adjustor on an Exchange 5.5 server, and instruct it to remove invalid users from public folder ACLs before the upgrade. Ideally, this should be done on one server, before the first Exchange 2000 installation. Wait until the changes replicate to all other servers. To start the Directory Service/Information Store consistency adjuster in the Exchange 5.5 Administrator program, right–click the server object, click **Properties,** click the **Advanced** tab, and then click **Consistency Adjuster.**

The second possibility can occur when a Windows 2000 user without a mailbox becomes the Exchange Administrator, creates a public folder, and sets the client permissions of the folder so that the default permission is Owner, leaving a Windows 2000 user on the ACL. Exchange *5.5* users cannot see the public folder because the Windows 2000 user without an Exchange *5.5* mailbox cannot be represented in an Exchange *5.5* server ACL.

Distribution Lists and Security Groups

In earlier versions of Exchange Server, you could assign permissions to public folders by using distribution lists. Exchange 2000 uses Windows 2000 universal security groups for the same purpose (universal security groups are associated with ACLs). However, unless your Windows 2000 forest contains at least one native–mode domain, universal security groups cannot exist.

Therefore, you can create at least one native–mode domain (known as a group management or transition domain) in a forest populated with distribution lists from your existing Exchange *5.5* organization. Initially these distribution lists are stored in Active Directory as universal distribution groups.

When you run an in-place upgrade, and Active Directory Connector has created universal distribution groups in a native–mode domain, public folders with distribution lists will have permissions replaced with universal distribution groups. The first time a member of a universal distribution group attempts to gain access to the contents of that folder, Web Storage System attempts to change the universal distribution group to a universal security group. If the universal distribution group was created in a mixed–mode domain, the upgrade to universal security group fails. An event will record which universal distribution group caused the upgrade process to fail, so that administrators can correct the problem. If the universal distribution group was from a native–mode domain, Web Storage System upgrades it to a universal security group.

The upgrade process can be controlled further. Sometimes, Windows 2000 administrators want to restrict the change of universal distribution groups to universal security groups. Control this by setting the *msExchDisableUDGConversion* attribute on the organization object. If the value of this attribute is 0 (the default), Web Storage System automatically converts universal distribution groups to universal security groups. If the value of this attribute is 1, Web Storage System does not automatically convert universal distribution groups. If the request to do so comes from a client, replication and upgrade can still convert universal distribution groups to universal security groups. If the value of this attribute is 2, Web Storage System does not convert a universal distribution group to a universal security group.

If you do not want to use universal security groups or switch Windows 2000 domains to native–mode, use domain local or global groups for public folder access control. However, this method requires continual maintenance and administration.

Public Folder Administration

Another change in Exchange 2000 is the public folder administration model. With earlier versions of Exchange, you can connect to a server and administer public folders on the server if you are the server's Exchange Full Administrator (secure folders are handled differently). If the folder in an earlier version of Exchange is not secure, users are Exchange *5.5* administrators when they have administrative rights on the public information store to which they are connecting, or when the users have administrative rights on an object higher in the directory service hierarchy. When logged on to the public information store, these administrators can perform administrator–only actions on all public folders in the hierarchy.

In Exchange 2000, public folder administration is based on permissions set on various objects (top level hierarchies, administrative groups, and the public folders themselves). In Exchange 2000, various permissions can be set for various public folder actions, such as adding replicas or changing age limits, rather than rights for all actions in Exchange *5.5.* Although Exchange 2000 is as consistent as possible with Exchange *5.5,* the administration model does not match exactly, creating some administrative issues.

A public folder tree is a hierarchical grouping of public folders that can exist on one or more Exchange 2000 public folder stores. In Exchange 2000, there can be multiple public folder trees, and each tree has its own public folder store. **A** public folder store can host only one public folder tree. In Exchange Server *5.5,* the system is constrained to a single public folder tree called Public Folders. In Exchange *5.5,* placeholders for all public folders within a public folder tree are represented on all public information stores associated with that public folder tree. Exchange 2000 supports the MAPI public folder tree similar to that of Exchange Server *5.5,* but supports multiple non-MAPI public folder stores.

ADC populates and synchronizes Active Directory with public folder information from the Exchange *5.5* directory. You can create connection agreements for both recipients (which synchronize mailboxes, custom recipients, and distribution lists) and public folders between a Windows 2000 server and an Exchange *5.5* server. You can only create connection agreements for public folders where an Exchange directory container exists in the domain with which you want to make the connection. When you type the name of the Exchange *5.5* server on the ADC connection agreement **Connections** tab, the values on the **From Windows** and **From Exchange** tabs are automatically entered into the Exchange directory container.

To move public folders to Exchange 2000, replicate existing public folders on the old Exchange server to the new Exchange 2000 server, and then remove the instances from the old Exchange server. Users then automatically connect to the Exchange 2000 public folder when they next attempt access.

When you upgrade an existing Exchange *5.5* server to Exchange 2000, all the existing Exchange sites match one-to-one with Exchange 2000 administrative groups and routing groups.

Public Folder Affinity and Referrals

If an Exchange 2000 server is installed into an existing Exchange 5.5 site that uses public folder affinity to access folders in remote sites, both Exchange 2000 users and Exchange 5.5 users can access the remote folders. By default, Exchange 2000 allows public folder referral over an Exchange 5.5 network connection.

In Exchange 2000, public folder affinity is transitive, therefore if an affinity exists between site 1 and site 2 and between site 2 and site 3, an affinity automatically exists between site 1 and site 3. In Exchange 5.5 and earlier versions, public folder affinity is not transitive. In a typical Exchange 2000 organization, routing groups are connected for efficient mail–flow. However, because public folder affinity is transitive, and if you allow public folder referrals, public folder referrals are then available to all other servers in the connected routing groups. You can avoid this by restricting public folder referrals in a connector's properties. Exchange administrators can select or clear the **Do not allow public folder referrals** check box on Exchange 5.5 site connectors and Exchange 2000 X.400 and SMTP connectors.

Circular Logging

In Exchange 2000, circular logging is disabled by default. This is different from earlier versions of Exchange because the use of transaction log files greatly enhances the recovery of an Exchange server. Databases within a storage group share transaction log files. Therefore, the circular logging setting applies to one storage group at a time. For more information about circular logging, see the Exchange Web site at http://www.microsoft.corn/exchange.

Communication Protocols

There is support for standard communication protocols in Exchange 2000. In most cases, there are no special issues to consider when you upgrade.

Protocol Servers

There are several protocol virtual servers:

- SMTP
- NNTP
- Internet Message Access Protocol version 4 (IMAP4)
- Post Office Protocol version 3 (POP3)
- Web Distributed Authoring and Versioning (WebDAV)

These protocol virtual servers can operate together to provide a single protocol service for scalability and redundancy, and can span multiple computers.

These protocol virtual servers appear to clients as a single protocol virtual server. A particular virtual server is selected either by means of round-robin DNS. by Network Load Balancing, or possibly by a hardware device.

Outlook Web Access

Outlook Web Access in Exchange 2000 is different from the version introduced with Exchange Server version 5.0. That version of Outlook Web Access used Active Server Pages (ASP) to retrieve and render data from the Exchange 5.x information store. In Exchange 2000, rendering takes place directly in Web Storage System. Because of the new Outlook Web Access architecture, it is no longer possible to customize ASP files to change the basic user interface. You can, however, reuse the standard Exchange 2000 Outlook Web Access components in existing applications.

If you have modified the ASP files for Outlook Web Access in Exchange 5.5, consider initially keeping both versions of Outlook Web Access, and then gradually migrating your applications so that you can use the new Exchange 2000 development features. Use an Exchange 5.5 front-end Outlook Web Access server with an Exchange 2000 back-end server. Users make Hypertext Transfer Protocol (HTTP) requests to the Exchange 5.5 server, which then checks the Exchange directory for the data. A connection is then made to the Exchange 2000 back-end server, where the data is retrieved.

> **Note** You cannot use an Exchange 2000 front-end Outlook Web Access server with an Exchange 5.5 back-end server.

NNTP

NNTP content remains within the public folder structure when you upgrade from Exchange 5.5. However, the connector does not upgrade. Record your configuration settings and recreate the NNTP connector manually, after the upgrade process.

> **Note** When you attempt to store Exchange 2000 newsgroups in an NNTP virtual server that uses a network share, you may notice that the newsgroups are not populated with data. This happens when you set up an administrative network share. Use a regular network share instead of an administrative share to prevent this.

IMAP4

Exchange 5.5 IMAP4 settings and configuration are carried over during an upgrade to Exchange 2000.

POP3

Exchange 5.5 POP3 settings and configuration are carried over during an upgrade to Exchange 2000.

LDAP

Windows 2000 now provides the LDAP service. If you have an existing client using LDAP for directory searches, modify the LDAP server to point to a global catalog server in the forest.

> **Note** If the LDAP port number (389 by default) has been changed on your Exchange 5.5 server, Exchange 2000 ForestPrep and Setup fail. The Exchange 2000 ADC requires a connection to the default port number during the install process.

Link Monitors

Link monitors are not upgraded and need to be re-configured after the upgrade process.

Server Monitors

If you have created server monitors that are located in an Exchange 5.5 site and that monitor a server due for an upgrade, remove those monitors. An in-place upgrade of a server configured with server monitors removes the monitors. If you need to maintain this service, reconfigure the monitors on another Exchange server prior to the upgrade.

For a newly upgraded Exchange 5.5 server or an Exchange 2000 server that has been installed into an existing Exchange 5.5 site, Exchange server monitors will be configured as follows:

- No notifications will exist.
- Status will show the server to be in a critical state if any of the following services are stopped:
  - Information store.
  - Message transfer agent (MTA) stacks.
  - Routing engine.
  - System Attendant.
  - SMTP.
  - World Wide Web publishing service.
- There will be no default warning state.

In all cases, earlier versions of Exchange should only monitor earlier versions of Exchange. Exchange 2000 servers should only monitor Exchange 2000 servers. An incorrect configuration can result in false alerts.

Message Tracking

In Exchange 5.5, message tracking can be turned on in three locations. In Exchange 2000, message tracking is on or off for the entire server. During an in-place upgrade, if message tracking is on in any of the locations on an Exchange 5.5 server, it will remain on for the entire server when the upgrade is complete.

Tracking Log Format

The tracking log format in Exchange 2000 differs significantly from the format in Exchange 5.5. If you are using third–party reporting tools to collect data from Exchange 5.5, or if you have monitoring software that tracks event logs, obtain updates from your vendor.

Exchange 2000 message tracking logs can improve the message identification: you can log the message subject in the message tracking log files. This new feature also enables the display of the message subject in the queue viewer in Exchange 2000. Message tracking logs are stored in Exchsrvr*servername*.log and contain information about senders, the time a message was sent or received, message size and priority, recipients, and message subjects if you have enabled subject logging.

Message Event Log Format

Message tracking logs troubleshoot the flow or status of the message in the Exchange 2000 system. The logs have a column labeled Event-ID that contains information about actions for a message such as sent, received, deleted, and retrieved.

Transport Stacks and Connectors

Special considerations for transport stacks and connectors are included in the following sections. There are no special considerations for the following:

- TCP/IP
- X.25
- MS Mail connector
- MS Mail directory synchronization connector
- Schedule+ Free/Busy connector
- Lotus cc:Mail connector
- Lotus Notes connector

Remote Access Service

Exchange 2000 does not support remote access service (RAS) X.400. You cannot upgrade an Exchange 5.5 Dynamic RAS connector to Exchange 2000. Instead, you must migrate your Exchange 5.5 Dynamic RAS connectors to Exchange 2000 SMTP connectors or TCP/IP RAS connectors.

TP4

Transport Class 4 (TP4) is no longer supported in Windows 2000. If connectivity is required using this protocol, retain a server running Exchange 5.5 or earlier on Windows NT Server 4.0.

X.400 Connectors

X.400 connectors are configured from one bridgehead server to another. When a server with an X.400 connector is upgraded to Exchange 2000, the setup process verifies that the remote MTAs for these connectors are on Exchange 2000 servers. If so, the setup process attempts to upgrade the X.400 site connector and the corresponding remote X.400 site connector to Exchange 2000 SMTP connectors. If the remote MTA is an Exchange 5.5 server, or if it is not possible to upgrade the remote MTA to SMTP (for example, permissions do not allow this) the upgraded server's X.400 connectors upgrade to Exchange 2000 X.400 routing group connectors. In addition, the stacks that the X.400 connectors are configured on are also upgraded to Exchange 2000 X.400 stacks.

Each X.400 Connector object must be associated with one of two transport stacks, X.25 (TP0) or TCP/IP.

From the directory service perspective, migration of an X.400 connector is uncomplicated. There are no delivery restrictions and message size limit settings are replaced with the new cost categories.

The following table shows the Exchange 5.5 X.400 connector properties (the X.400–link object class) and how they correspond to the Exchange 2000 X.400 connector properties.

Table 10.2 Exchange 5.5 X.400 attributes and Exchange 2000 equivalents

Description	Exchange 5.5 Attribute	Exchange 2000 Attribute	Comments
Display name	Admin-Display-Name	Admin-Display-Name	
Directory name	Common–Name	Common–Name	
Cost	Part of address space and connected sites properties	Normal–Cost	New to Exchange 2000
Local MTA	Home–MTA	Home–MTA	
Remote MTA	Gateway-Local-Desig	Gateway-Local-Desig	
Schedule	Activation-Style, Activation-Schedule	Activation-Style, Activation-Schedule	Style refers to whether to always run, use schedule, or never run
Connected sites	Connected–Domains	None	One connected routing group per Exchange 2000 X.400 connector

Table 10.2 Exchange 5.5 X.400 attributes and Exchange 2000 equivalents *(continued)*

Description	Exchange 5.5 Attribute	Exchange 2000 Attribute	Comments
Address space	Routing–List	Routing–List (Only on external gateway X.400 connectors)	
Delivery restrictions	Auth-Orig, Unauth-Orig	VIP-Cost	New to Exchange 2000
Message size	Deliv-Cont-Length	Large–Cost	New to Exchange 2000
Stack	P-Selector, P-Selector-Inbound, S-Selector, S-Selector-Inbound, T-Selector, T-Selector-Inbound	P-Selector, P-Selector-Inbound, S-Selector, S-Selector-Inbound, T-Selector	Remote stack selectors. Some MTAs have different inbound and outbound selectors. T-Selector-Inbound is eliminated for Exchange 5.5 SP1 and Exchange 2000.
Override	MTA-Local-Cred, MTA-Local-Desig	MTA-Local-Cred, MTA-Local-Desig	Override name and password for local MTA.
Maximum open retries	Num-Of-Open-Retries	Num-Of-Open-Retries	Number of times the MTA tries to establish a remote connection (default value is 20).
Maximum transfer retries	Num-Of-Transfer-Retries	Num-Of-Transfer-Retries	Number of times the MTA tries to resend a block of data (default value is 20).
Open interval	Open-Retry-Interval	Open-Retry-Interval	Number of seconds the MTA will wait before trying to re-connect (default value is 60).
Transfer interval	Transfer-Retry-Interval	Transfer-Retry-Interval	
Checkpoint size	RTS-Checkpoint-Size	RTS-Checkpoint-Size	Size of data in kilobytes that is transferred in a block (1–4).
Recovery timeout	RTS-Recovery-Timeout	RTS-Recovery-Timeout	Maximum number of seconds before aborting recovery of a connection (default value is 10).

Table 10.2 Exchange 5.5 X.400 attributes and Exchange 2000 equivalents *(continued)*

Description	Exchange 5.5 Attribute	Exchange 2000 Attribute	Comments
Window size	RTS-Window-Size	RTS-Window-Size	Number of checkpoints sent before an acknowledgement (ACK) response is expected.
Association lifetime	Association-Lifetime	Association-Lifetime	Number of seconds an association is idle before it is removed (default value is 60).
Association disconnect	Session-Disconnect-Timer	Session-Disconnect-Timer	Number of seconds the MTA waits for an acknowledgement (ACK) response after a session disconnects.
Association threshold	Temp-Assoc-Threshold	Temp-Assoc-Threshold	When the number of queued messages reaches this value, another association is started.
Transfer timeouts	Transfer-Timeout-Non-Urgent, Transfer-Timeout-Urgent, Transfer-Timeout-Normal	Transfer-Timeout-Non-Urgent, Transfer-Timeout-Urgent, Transfer-Timeout-Normal	
Message text word wrap	Line–Wrap	Line–Wrap	
Remote clients support MAPI	Encapsulation–Method	Encapsulation–Method	
MTA conformance			1984, 1988 X.410, 1988 normal
X.400 link options	Two-Way-Alternate-Facility	Two-Way-Alternate-Facility	
Global domain identifier override			Overrides the default global domain identifier that is constructed from the local X.400 address space

Note When you create an address space on an Exchange 2000 X.400 connector without an entry in the Administrative Management Domain Name field, the Administrative Management Domain Name field receives no value. The Administrative Management Domain Name field is a mandatory X.400 field. In earlier versions of Exchange, this field has a default value of a blank space. This is by design. Although the Administrative Management Domain Name field is a mandatory field for X.400 addresses, it is not necessary that it be expressly filled for address space entries on connectors. A field with no value is considered a wildcard and matches any value.

Site Connectors

Site Connectors are configured as one-to-many or all-to-many connectors. The Site Connector is X.400, remote procedure call (RPC)–based, and easy to configure in Exchange 5.5. In Exchange 2000, SMTP connector configuration is simple, and within an organization, routing group connectors are also simple to configure.

When a server in an Exchange 5.5 site is upgraded to Exchange 2000, and if that server is the local MTA for a one-to-many Site Connector, or if there are all-to-many Site Connectors in the site, Exchange Setup verifies which of the target servers of these connectors are Exchange 2000 servers. That is, this check determines whether at least one Exchange 2000 server on both sides of any existing Site Connectors exists within the local site. If so, Setup attempts to upgrade the local and remote sides of these Site Connectors to Exchange 2000 SMTP inter-routing-group connectors. If this is not possible (for example, if permissions do not allow this), the local Site Connectors are upgraded to Exchange 2000 X.400 RPC inter-routing-group connectors. The Exchange 2000 X.400 RPC connector can only be used when connecting to an Exchange 5.5 Site Connector. Routing groups with only Exchange 2000 servers cannot use the Exchange 2000 X.400 RPC connector for inter-routing-group connectivity.

Site Connectors can only be used in Exchange 2000 to link to Exchange 5.5 sites. Remove the Site Connector from an environment with only Exchange 2000 servers. Use an SMTP inter-routing-group connector instead. Consider configuring a Site Connector to send messages from selected servers in the site instead of configuring the Site Connector for either one or all servers. This is possible after an SMTP connector replaces the Site Connector.

For connections from Exchange 2000 to Exchange 5.5 sites, the properties of the Site Connector are carried over during upgrade with the exception of Cost, which becomes multiple attributes: Normal-Cost, Urgent-Cost, VIP-Cost, and Large-Cost.

The following table shows the properties for the Site Connector.

Table 10.3 Exchange 5.5 Site Connector attributes and Exchange 2000 equivalents

Description	Exchange 5.5 Attribute	Exchange 2000 Attribute	Comments
Display name	Admin-Display-Name	Admin-Display-Name	
Directory name	Common-Name	Common-Name	
Local site	Extracted from Obj-Dist-Name	None	
Target site	Domain-Name	Domain-Name	
Cost (overall connector)	Cost	Large-Cost, VIP-Cost, Normal-Cost, Urgent-Cost	New cost categories for Exchange 2000
Costs (per target server)	Target-MTAs	None	
Target servers	Target-MTAs	Target-MTAs	
Address space	Routing-List	None	Address spaces are not used on inter-routing-group Exchange 2000 connectors
Local bridgehead server	Home-MTA (if non-existent, all local site servers are bridgehead servers)	Home-MTA (if non-existent, all local routing group servers are bridgehead servers)	
Override account	Authorized-Domain, Authorized-User, Authorized-Password	Authorized-Domain, Authorized-User, Authorized-Password	

Dynamic RAS Connector

Exchange 2000 does not support Dynamic RAS Connectors. You cannot upgrade an Exchange 5.5 Dynamic RAS Connector to Exchange 2000. Instead, migrate your Exchange 5.5 Dynamic RAS Connectors to Exchange 2000 SMTP connectors, or TCP/IP RAS Connectors.

Internet Mail Connector

When an Exchange 5.5 server running an Internet Mail Connector is upgraded to Exchange 2000. Setup replacer the Internet Mail Connector with an Exchange 2000 SMTP connector.

The lnternct Mail Connector is replaced by the new Windows 2000 SMTP service. The SMTP service adds an extra dimension to the capabilities of the former Internet Mail Connector by adding virtual servers. With the Internet Mail Connector, only one logical and physical instance of SMTP can run on a server. Virtual servers are logical instances of SMTP with unique IP addrerslport assignment cornbinations and fully qualified domain names.

When an Exchange 5.5 server is upgraded to Exchange 2000, an Internet Mail Connector is replaced with a single SMTP virtual server.

Table 10.4 shows how internet Mail Connector properties are upgraded in Exchange 2000:

Table 10.4 Exchange 5.5 and Exchange 2000 Internet Mail Connector features

Exchange 5.5 Feature	Upgraded	Comments
Internet Mail Tab		
Administrator's mailbox	No	
Notifications	No	Default notification rules for Exchange 2000 are used.
Message content	Yes	Default message content configuration maps to default Internet message format on Exchange 2000.
Character sets	Yes	Default character set configuration maps to default Internet message format on Exchange 2000.
E-mail domain	Yes	Per-domain content settings maps to multiple Internet message formats on Exchange 2000, listed by address space.
Client Support S/MIME Options	No	
Convert inbound message to fixed–width font	No	
Enable Message Tracking	No	
Advanced Options	Partial	Out-of-Office and Auto–reply values map to Exchange 2000 Internet message formats.
Connections Tab		
Transfer Modes	No	

Table 10.4 Exchange 5.5 and Exchange 2000 Internet Mail Connector features *(continued)*

Exchange 5.5 Feature	Upgraded	Comments
Transfer Modes Advanced button	Partial	Maximum number of inbound and outbound messages map to the Exchange 2000 default SMTP virtual server.
Use DNS	Yes	
Forward all messages to host	Yes	
Dial using	No	
Specify by e-mail domain	No	
Accept connections from any host	Partial	Defaults to allow use of all connections. Administrators need to manually modify the virtual server after upgrade.
Only from hosts using	Partial	Defaults to allow use of all connections. Administrators need to manually modify the virtual server after upgrade.
Specify by host button	Partial	Based on first entry read, either the granted access list or the denied access list moves to the default SMTP virtual server **Access** tab or **Connections** button. Entries not mapped are listed in the progress log.
Message filtering button	Partial	All filtered domains will map to Global Settings\Message Delivery properties\Filtering tab with an '@' mapped to all domains that did not have filtering in Exchange 5.5. The **Enable filtering** switch is selected for the port in the default SMTP virtual server. **Delete messages instead of moving to the Turf directory** is not mapped.
Retry interval (hrs)	Partial	Only the first four intervals are mapped and these convert from hours into minutes.
Time–outs button	Partial	No distinction is made for time-outs based on priority. Normal message timeout and usual delay notification settings map to the default SMTP virtual server.

Table 10.4 Exchange 5.5 and Exchange 2000 Internet Mail Connector features *[continued]*

Exchange 5.5 Feature	Upgraded	Comments
Clients can only submit if the user's mailbox is on this server	No	
Clients can only submit if authentication account matches submission address	No	
Queues Tab		
Queued mail inbound	Yes	Mail in the Internet Mail Connector's inbound queue gets moved to Exchange 2000 \exchsrvr\mailroot\vsi #\pickup directory.
Queued mail outbound	Yes	Mail in the Internet Mail Connector's outbound queue will get moved to Exchange 2000 \exchsrvr\mailroot\vsi #\pickup directory.
Mail queued for ETRN	No	If you configure a domain on the Internet Mail Connector to hold messages for remote de-queuing (ETRN or TURN), after upgrade to Exchange 2000, messages queued for ETRN or TURN get moved to the Exchange 2000 queue and return a non-delivery report (NDR) immediately.
Routing Tab		
Do not re-route incoming SMTP mail	Yes	Maps to Exchange 2000 default SMTP virtual server properties, **Access** tab, **Relay Restrictions** button.
Re-route incoming SMTP mail	Yes	Maps to Exchange 2000 default SMTP virtual server properties, **Access** tab, **Relay Restrictions** button.
Inbound routing table	No	
Security Tab		
Per-domain **Security** property page	No	
General Tab		
Computer name	Yes	Maps to the bridgehead server on the Exchange 2000 SMTP connector.

Table 10.4 Exchange 5.5 and Exchange 2000 Internet Mail Connector features *(continued)*

Exchange 5.5 Feature	Upgraded	Comments
Message size	Yes	Maps to the **Content Restrictions** tab, **Allowed Sizes** option on the Exchange 2000 SMTP connector.
Administrative note	No	
Connected Sites Tab		
Connected sites table	Yes	All connected sites with organization name, site name, routing address, and cost map to the connected routing groups table in Exchange 2000.
Address Space Tab		
Address space table	Yes	All address spaces map to the Exchange 2000 SMTP connector with Type, Address, and Cost preserved. The scope changes to the most restrictive type for all address rpaces after upgrade.
Delivery Restrictions Tab		
Accept messages from table	Yes	Maps to Exchange 2000 SMTP connector's **Delivery Restrictions** tab.
Reject messages from table	Yes	Maps to Exchange 2000 SMTP connector's **Delivery Restrictions** tab.
Diagnostic Logging Tab		
Diagnostic logging table	No	
Registry Entries		
MaxRecipients	Yes	Maps to default SMTP virtual server properties, Messages tab, **Limit number of recipients per message to** field.
MaxReceivedHeaders	Yes	Maps to default SMTP virtual server properties, **Delivery** tab, **Advanced** button, **Maximum hop count** field.
DisableOutboundESMTP	No	
DisableReverseResolve	No	

Table 10.4 Exchange 5.5 and Exchange 2000 internet Mail Connector features *(continued)*

Exchange 5.5 Feature	Upgraded	Comments
DefaultDomain	Yes	Maps to default SMTP virtual server properties, **Delivery** tab, **Advanced** button, **Fully–qualified domain name** field.
AppendDefaultDomain	Yes	Maps to default SMTP virtual server properties, **Delivery** tab, **Advanced** button, **Masquerade domain** field.

PROFS Connector

The Professional Office System (PROFS) or Office Vision connector is not supported in Exchange 2000. If continued connectivity is required, retain an Exchange 5.5 server in a connected site.

SNADS Connector

The SNA Distribution System (SNADS) connector is not supported in Exchange 2000. If continued connectivity is required, retain an Exchange 5.5 server in a connected site.

Other Exchange Services

The following Exchange services require special planning when you upgrade to Exchange 2000.

Chat Services

Use the chat migration tool to create a new chat community on an Exchange 2000 server with the name of the Exchange 5.5 chat server, and with most of the same settings and configurations. The chat migration tool, Chatmig.exe, is located in the Exchsrvr\Bin directory. This tool will not migrate Chat Service 5.5 extension data. If you have a Profanity Filter or Transcription Extension with your Chat Service 5.5, you must migrate these manually.

The chat community you create with the migration tool is separate from your Exchange 5.5 chat server. Exchange 2000 Server does not support an Exchange 5.5 chat installation that consists of one chat community that spans a network of two or more Exchange servers. Use the migration tool to create two similar chat communities on each Exchange 2000 server. The chat communities on the Exchange 2000 servers in your domain operate separately, with separate user lists and channel lists.

For step-by-step instructions for the chat migration tool, see the Exchange 2000 Server online documentation.

Key Management Service

In-place upgrade is supported if the Key Management Service and Certificate Services are on same operating system. The underlying code is upgraded during Exchange Setup. When the Key Management Service starts, the original Key Management Service database is upgraded.

In Exchange 5.5, when users are deleted, their key history and certificates remain in the Key Management Service database. In Exchange 2000, when users are removed from Active Directory, their key history is removed.

Do not install the Exchange 2000 version of Key Management Service into a mixed-mode organization. Instead, continue using the existing Exchange 5.5 Key Management Service and manage it from the Exchange 5.5 Administrator program.

Event Scripts

Exchange 2000 runs an Exchange 5.5 compatible event-script service. Event scripts run with no modification. Because security is stronger in Exchange 2000, some permissions in the scripts may need to be changed. Because event scripting in Exchange 2000 is more powerful, you need to rewrite scripts to take advantage of Exchange 2000 functionality.

Third-Party Software

It is unlikely that your existing third-party tools, connectors, virus checkers, and so on will work correctly after you upgrade an Exchange 5.5 server to Exchange 2000. Most third-party applications use the Directory application programming interface (DAPI) protocol to access the Exchange 5.5 directory. This is not supported within Active Directory. Applications must be rewritten to use ADSI or LDAP to manipulate Active Directory. You should either disable or remove third-party software prior to upgrade.

Client Effects

When preparing for your Exchange 2000 deployment it is important to understand how it will affect the user interface for message clients within the organization. This section looks at each type of client and the impact of an upgrade or migration.

Exchange Client, Outlook 97 and Outlook 98

Active Directory replaces the local directory service on earlier versions of Exchange. However, earlier clients are not aware of Active Directory and continue to direct MAPI requests to the local Exchange server for address searches, or attempt to resolve names in message address fields. Exchange 2000 cannot respond directly and sends the requests to a global catalog server. This process is carried out by the DSProxy service. DSProxy uses DNS to locate a local global catalog server in a Windows 2000 site to service MAPI requests.

Outlook 2000

Outlook 2000 manages MAPI requests differently. Like earlier Outlook clients, the first request is targeted to the user's Exchange 2000 server. The server's DSProxy service, instead of forwarding the request, responds with a referral pointing the client directly to the local global catalog.

Note Do not attempt to install Outlook 2000 and Microsoft Exchange 2000 Server on the same computer. This is not supported by design.

Outlook Profiles

For clients with Outlook 2000 and earlier versions of Outlook, upgrading to Exchange 2000 has no effect on client profiles, logging on, or client behavior. If you deploy Exchange 2000 into a mixed-mode organization and move users, profiles are updated to reflect the change provided that you maintain the old server in the administrative group. This behavior is identical to moving users between two Exchange 5.5 servers within the same site.

POP3 and IMAP4 Clients

The deployment of Exchange 2000 front-end and back-end servers affects what needs to be reconfigured in your POP3 and IMAP4 clients.

For example, assume an IMAP4 or POP3 client is configured to point to one Exchange 5.5 server for its incoming mail, outgoing (SMTP) mail, and LDAP directory service. If you upgrade your Exchange 5.5 server, the incoming and outgoing server fields remain valid, but the LDAP server needs to be changed if your Exchange 2000 server is not a global catalog server.

A front-end and back-end architecture simplifies your options. Mailboxes can be moved to other Exchange 2000 servers in the same site. The original Exchange 5.5 server can be upgraded to Exchange 2000 and configured as a front-end server. POP3 and IMAP4 requests can then go to the back-end mailbox servers.

PART 4

Basic Deployment Planning

In This Part

Administration and Maintenance

Server Design for Backup and Restore

Virus Protection

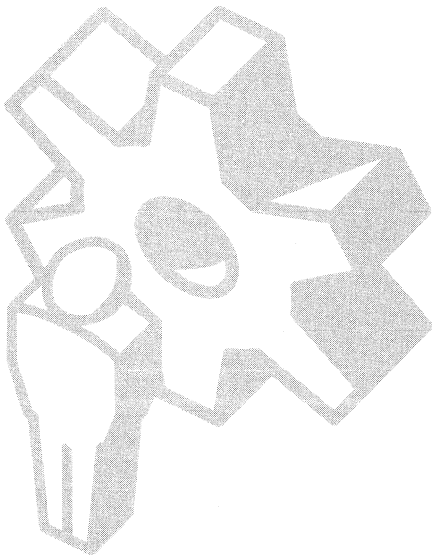
Server Availability

Server Sizing

Message Routing

Backbone Configuration and Tuning

External Connectivity



CHAPTER 11

Administration and Maintenance

Robert Dring, Senior Consultant, Microsoft
Stanley Lum, Consultant, Microsoft
Jens Trier Rasmussen, Senior Consultant, Microsoft

In the past, many organizations divided their information services staff into multiple groups. One group was responsible for handling all aspects of the messaging system, including server and recipient management. Another group was in charge of handling Microsoft Windows NT account administration. A third group handled the management of Windows NT servers. With the introduction of Microsoft Windows 2000 Server and Microsoft Exchange 2000 Server, these divisions must be re-evaluated.

Exchange 2000 has some new features that require re-evaluating existing administrative models. Recipient administration is now unified with Windows 2000 account administration and is handled using the Active Directory Users and Computers snap-in in Microsoft Management Console (MMC). Server administration is handled by using the Exchange System Manager MMC snap-in; it has been divided into administrative and routing groups.

In Exchange 2000, the same administrative group can handle recipient and server management. It is also possible to centrally manage all Exchange servers, while delegating recipient management to business units.

This section discusses several new ways to administer recipients and servers and highlights key monitoring tools available to ensure smooth system operation.

In This Chapter

Recipient Management

Server Management

Additional Administrative Features

Mixed Mode Operation with Exchange 5.5

Managing and Monitoring Exchange Server Performance

Recipient Management

Exchange 2000 Server has been redesigned to reflect more accurately how most customers use Exchange in their network environments. Exchange allows you to use an administrative model to organize the Exchange 2000 Server topology so it meets your business needs.

With Exchange 2000, recipient management is now a matter of manipulating object attributes in the Windows 2000 directory service. called Active Directory. When an Active Directory user object is mail-enabled, Exchange-specific attributes, such as mailbox server and e-mail addresses, are populated to that particular object. Changing a user's phone number or Simple Mail Transfer Protocol (SMTP) address requires connecting to a domain controller and modifying the attributes of a user object in Active Directory. Contrast this to Exchange version 5.5, where changing a user's phone number or SMTP address requires connecting to an Exchange server and changing the attributes of a mailbox object in its directory.

If your organization includes servers running Exchange 5.5, it is important that you understand how Exchange objects (from previous versions of Exchange) have changed in Exchange 2000. Table 11.1 shows how Exchange 5.5 concepts are represented in Exchange 2000.

Table 11.1 Exchange 2000 terminology

Exchange 5.5 Concept	Exchange 2000 Equivalent
Mailbox	Active Directory mailbox-enabled user
Custom recipient	Active Directory mail-enabled contact or user
Distribution list	Active Directory mail-enabled group
Public folder (visible in directory)	Active Directory public folder

Organizational Units

Before selecting an administrative model, you must have a good understanding of organizational units. Objects in Active Directory exist within domains. They can be further organized within a domain by using organizational units. Permissions can be set on domains and organizational units to delegate administrative authority over the contents. Delegating control by using organizational units is the preferred method, because you can easily move objects in a domain between organizational units. Creating extra domains requires additional domain controllers and makes moving objects more difficult. Although this method is appropriate when security needs are different in each domain, it is probably inappropriate for administrative delegation. Note that mail-enabled user objects do not have to be located in the same organizational unit or even in the same domain as the Exchange 2000 server the mailbox physically resides on.

In Exchange 5.5, creating multiple recipient containers for mailboxes is not recommended because moving mailboxes between containers is difficult. With Exchange 2000 and Active Directory, moving an object from one organizational unit to another is a simple operation that does not affect the associated mailbox. You can easily create an organizational unit structure to delegate administrative control over user objects. Moving objects between domains is a more complicated task that requires additional tools, such as MoveTree or Active Directory Migration Tool. However, even moving a user object between domains does not affect the associated Exchange 2000 mailbox. For more information on MoveTree and Active Directory Migration Tool, see the Windows 2000 documentation.

Mailboxes can be moved between mailbox stores on an Exchange server and between servers, whether or not the user objects are moved in Active Directory. Exchange 2000 also allows you to move mailboxes between administrative groups when Exchange 2000 is operating in native mode. However. it is not possible to move a server from one administrative group to another.

Administrative Models

You can use an administrative model to organize your Exchange 2000 Server system topology in a manner that addresses your business needs and process-related requirements. There are three main administrative models for managing Exchange 2000 recipients:

- Centralized administration
- Delegated administration
- Messaging group administration of Exchange attributes

Windows 2000 uses the centralized and delegated models. The third model—messaging group administration of Exchange attributes—is suitable for organizations that want to maintain the organizational structure used by Exchange 5.5 and Windows NT 4.0. Windows 2000 Active Directory and Exchange 2000 provide the flexibility to allow you to implement any of these administrative models.

When you select an administrative model, keep in mind that you should use the centralized or delegated administrative models and match them to the storage group structure of the Exchange organization. In either case, it is recommended that you combine all aspects of user administration and that you have the same people handle both account and mailbox issues.

Centralized Administration

To implement a centralized administration model, you must group user objects into organizational units. This allows you to control the application of Group Policy. The centralized account management group should have administrative rights on all the organizational units.

Delegated Administration

To implement a decentralized, delegated administrative model, administrative rights for individual organizational units are delegated to the appropriate groups. The organizational unit structure takes into account both the application of Group Policy and the delegated administrative model.

Messaging Group Administration of Exchange Attributes

Mailboxes are not separate objects in Active Directory, but are attributes of the user object. It is possible to separate administration of Exchange attributes from other attributes on user objects by specifying attribute-level permissions. However, this results in an unwieldy administrative model, which loses many of the benefits of directory integration.

g dminis Tools and tures for Recipie
m nt

After you have chosen an administrative model, you need to implement it. Before you do so, you need to understand the Exchange 2000 and Windows 2000 features that are important to recipient management. This section discusses these features in detail.

Administrative Tools

To manage computers running Windows 2000 Server remotely, you can install Windows 2000 Administrative Tools on Windows 2000 Professional clients; this provides the MMC snap-ins and other necessary tools. To install these tools from the Windows 2000 Server CD, open the I386 folder, and then double-click the Adminpak.msi file.

You can install the tools to manage Exchange 2000 remotely on Windows 2000 Professional clients by selecting Exchange 2000 System Management Tools in the Exchange 2000 Setup wizard.

An alternate method for installing these tools is to use the software installation component of the Group Policy snap-in on a server running Windows 2000 Server. You can use this tool to assign particular programs to selected computers or to publish them in Active Directory; this way administrators can use Add/Remove Programs in Control Panel to install the tools.

Windows 2000 Terminal Services can be installed on any server that is operating in remote administration mode. This allows users with appropriate permissions to remotely access a console session on a server over a LAN or dial-up connection. This capability is provided automatically with Windows 2000 Server and provides the following advantages to administrators:

- Control over all aspects of the server without the need to be physically present at the server.
- Ability to access the server from any client on which Terminal Services Client is installed.

  Administration capability over dial-up connections.

Exchange administrators should run Windows 2000 Professional on their workstations and install both the Windows 2000 Administration Tools and Exchange 2000 System Management Tools on them.

In addition, Windows 2000 Terminal Services should be installed on all servers that operate in remote administration mode. Then, an administrator whose workstation runs any version of Windows and the Terminal Services Client can control the servers from anywhere.

Organizational Units

Unlike recipient containers in Exchange 5.5, organizational units do not show up in the Microsoft Outlook Address Book and do not affect the ability to move mailboxes. With Exchange 5.5, it is recommended that administrators avoid creating recipient containers for administrative delegation, because doing so makes it difficult to move mailboxes between containers. With Exchange 2000, it is appropriate to create organizational units as needed for delegation of user object administration and the application of Group Policy, because objects can easily be moved between any of the organizational units in a domain.

Create the organizational unit structure that supports administrative delegation and application of Group Policy. The structure that you choose does not affect your ability to move mailboxes in Exchange 2000 or the ability of Outlook clients to view their address books.

Active Directory Delegation of Control Wizard

Use Active Directory Delegation of Control Wizard to grant access rights to domains or organizational units. This tool allows for fine control over delegation of authority. For example, you can allow or deny the ability to read or write particular properties of certain object classes. For more information, see the Windows 2000 documentation.

Keep things simple wherever possible by delegating authority over entire organizational units. Avoid specifying different permissions for particular properties, because they can affect other organizational units.

Exchange Administration Delegation Wizard

You use the Exchange Administration Delegation Wizard to grant rights to managing and configuring Exchange 2000 servers. Administration delegation is discussed in more detail in "Server Management," later in this chapter. You can use the wizard to assign the administrative group view role over particular Exchange administrative groups. Only users who have Exchange Administrator rights over one or more administrative groups can create mailboxes by using the Exchange extension to the Active Directory Users and Computers snap-in.

Use the Exchange Administration Delegation Wizard to create groups for the purpose of delegating administration, rather than specifying individual accounts. This makes it easy for you to change who has administrative rights by simply changing group membership.

Permissions

Users with accounts in the Account Operators, Administrators, or Domain Administrators groups have full control over user objects and can read and modify all attributes, including mailbox location and other Exchange-related information. Active Directory Users and Computers imposes a further restriction on the ability to make a user mailbox-enabled or to move mailboxes: you can create mailboxes in the Exchange administrative groups only if you have Mailbox Manager rights to those groups. If you do not have rights on any Exchange administrative groups, you cannot use Active Directory Users and Computers to make a user mailbox-enabled.

> **Important** This restriction applies only to the user interface. An administrator with appropriate rights to the user objects can modify the Exchange attributes by using the Active Directory Administration Tool (Ldp.exe) or another directory manipulation tool. For more information on the Active Directory Administration Tool, see the ***Windows*** 2000 Server *Resource Kit*.

By default, users in the Account Operators group have full control over user objects but no control over contacts. You can use the Active Directory Delegation Wizard to grant full control to the Account Operators group or to other groups. For more information about Active Directory Delegation Wizard, see the Windows 2000 documentation.

Use Exchange Administration Delegation Wizard to grant Mailbox Manager rights to those users who should have the ability to create mailboxes. Use the Active Directory Delegation Wizard to grant control over contacts to the appropriate groups.

Also keep in mind that, although it is possible to control permissions on individual user object attributes, administration will be greatly simplified by combining user account and mailbox administration roles into a single administrative group.

Recipient Policies

In Exchange 2000, recipient policies control c-mail address generation, much like the Exchange 5.5 site addressing feature. Recipient policies are more flexible than site addressing, because you can create multiple addresses of a given type and you can use filter rules to control which recipients the policies apply to. For example, an organization formed by the merger of two companies could use recipient policies to provide all users with an SMTP address reflecting the merged company name and different SMTP addresses reflecting the previous company affiliation.

Use recipient policies to assign multiple SMTP addresses to users that reflect every instance of a user's e-mail address, for example, it a user's name changes Assigning multiple SMTP addresses in this manner can eliminate the need to maintain alias filer on other systems.

Server Management

Issues that arise from Exchange 2000 Server management can be broken into three areas. It is important to understand these areas before you select an administrative model. Depending on the administrative model that you select, different groups in your organization need to focus on different management areas. The three main management areas are:

- **Exchange 2000** Administration and operation of Exchange servers using Exchange System Manager in MMC.
- **Windows 2000** Operating system configuration and management of the servers Exchange runs on.
- **Active Directory** Management and operation of the domain controllers that provide Active Directory services.

The same group of administrators can administer all three areas or a different group can administer each area separately. With either approach, responsibility can be centralized or delegated throughout the organization.

Managing Server Responsibilities

In the past, many corporations divided responsibility for Exchange servers between a messaging group that focuses on Exchange Server and a network operating system group that focuses on the underlying operating system. With the advent of Active Directory technology, corporations may consider splitting Active Directory operations from the operating system in the same way. This split works well when Exchange uses few servers and the application administrators do not want to manage issues involved with configuring the operating system. However, Exchange 2000 is typically a large-scale system, and may be the first PC-based application deployed across a corporation. This type of division can lead to communication problems between groups.

Typically, the members of the network operating system group are responsible for many different types of servers and have limited knowledge of Exchange. For example, Exchange 2000 has different backup and operating system requirements from most other Windows 2000 server applications. The messaging group has a clearly defined goal: keep mail flowing according to defined service level agreements. The network operating system group's goals are different. Whereas messaging administrators tend to think of their messaging servers as part of an enterprise system, network operating system group members typically treat each server independently as needed.

In a large-scale deployment of Exchange 2000, there could be up to hundreds of Exchange servers and dozens of Active Directory domain controllers. In addition, such organizations are likely to have hundreds of servers performing other tasks. Given the number of servers in each group, it is recommended that the administration of these three groups be divided and that each administrative group be given full responsibility for all aspects of each group of servers. This division of administration is shown in Table 11.2.

Table 11.2 Group responsibilities

Group	Responsibility
Messaging	Exchange 2000 and Windows 2000 configuration for all messaging servers.
Active Directory	Windows 2000 configuration for all Active Directory servers.
Network operating systems	Windows 2000 configuration for file and print servers, and other types of servers. Third-level support for the messaging and Active Directory groups. Other operating systems.

Administrative Models

You can manage Exchange 2000 servers centrally, independently of physical location. Alternatively, you can delegate management to multiple groups. Three administrative models are described in this section:

- Centralized management
- Distributed management
- Mixed management

When deciding on an administrative model, keep in mind that the centralized model is the simplest to create and manage. However, Exchange 2000 has the flexibility to support any administrative model that is dictated by your organization's structure.

Centralized Management

In a centralized management model, Exchange servers are divided into several administrative groups. Keep in mind that servers cannot be moved between administrative groups once they are placed within a group. In addition, there may be many routing groups to manage message transfer. This model is most appropriate for small to medium sized organizations that have a single information technology group, or in large organizations with a central information technology authority.

By default, servers are not organized into administrative groups in Exchange System Manager, because companies with a small number of servers do not usually have a need for administrative group divisions. For this reason, administrative groups are not displayed. However, any server that exists in your organization before displaying administrative groups will belong to the first administrative group. If you add servers to your organization before creating the administrative groups, the new servers will all belong to the first administrative group. If you want to add new servers to a new administrative group, you must first manually display administrative groups. Then you must create the new administrative groups before installing Exchange 2000 on the new servers. After servers are placed in an administrative group, they cannot be moved. If you do not want all servers to belong to the same administrative group, install Exchange 2000 on one server, create your administrative groups, and then install Exchange 2000 on the remaining servers and place them in the appropriate administrative groups.

Choosing properties on the Organization node in Exchange System Manager creates the first administrative group with all your servers; then you can define other administrative groups.

> **Caution** After you have defined the administrative groups, if you clear **Display Administrative Groups,** all servers are grouped into the same Servers node in Exchange System Manager and all group division is lost. If you then choose to display administrative groups again, all servers now belong to the first administrative group and are not restored to any prior group configuration.

> **Note** If your Exchange 2000 organization coexists with an Exchange 5.5 organization, administrative groups are always enabled.

Distributed Management

In a distributed management model, you divide Exchange servers into administrative groups by business unit or region, and you delegate administration responsibilities to the appropriate personnel. Within each administrative group, you assign servers to routing groups.

> **Note** When Exchange 2000 coexists with Exchange 5.5, an administrative group in Exchange 2000 appears and behaves like a site in Exchange 5.5.

Mixed Management

When an Exchange 2000 organization contains only Exchange 2000 servers, you can implement a mixed management model. With this model, you can create administrative groups to manage message routing, and you can create other administrative groups to apply server policies or to grant user rights for particular servers.

Understanding Administrative Tools and Features for Server Management

After you have chosen an administrative model, you need to implement it. Before you can do so. you need to understand the Exchange 2000 Server and Windows 2000 Server features that are important to server administration. This section discusser these features.

Domains

With Exchange 5.5, messaging servers are commonly grouped into resource doinains so messaging administrators can have domain administrator rights over those servers. Because of their greater flexibility, Windows 2000 organizational units are the best way to delegate such administrative rights. You do not need to create separate dornains in Active Directory to match Windows NT 4.0 resource domain structure.

Do not create domains solely for the purpose of delegating server administration rights. Use organizational units to delegate server administration rights.

Organizational Units

You use organizational units to group server objects for delegating administration. It is best to create an organizational unit structure for Exchange that reflects how you delegate Windows 2000 administration. When creating this structure, consider which groups should have the ability to start and stop services or to install new software.

Create an organizational unit to group Exchange 2000 rervers, and then grant access rights to the organizational unit to the Exchange administrators. If you implement a delegated administrative model, you can create additional organizational units.

Administrative Groups

Exchange servers can be grouped into administrative groups to delegate administration responsibilities. In addition, you can use administrative groups to separate management of system policies, routing groups, public folders, or chat communities. A server can belong to one administrative group for storage management purposes, and can also belong to a routing group in a different administrative group. You can apply system policies from another administrative group to all Exchange servers. However, you cannot create separate administrative groups for control over routing groups until the Exchange 2000 organization is in native mode.

When Exchange 2000 coexists with Exchange 5.5 in mixed mode, you should carefully plan the structure of the administrative groups that are used to define the administrative topology. Once you place a Microsoft Exchange 5.5 server or Exchange 2000 server in an Exchange 2000 administrative group, you cannot move that server to another administrative group. Your only option is to move an Exchange 5.5 server from an Exchange 5.5 only site to a mixed-mode Exchange 3000 organization. Even though an Exchange 2000 server cannot be moved once it is placed in an administrative group, you can move the data on the server. By using the Exchange Tasks Wizard. you car, move users' mailboxes between Exchange 2000 servers.

When Exchange 2000 is in native mode, you create an administrative group structure to match your chosen administrative model. If you have a delegated model, consider creating additional administrative groups to centralize some aspects of management, such as system policies. routing groups, public folders, and chat communities.

Routing Groups

You can create routing group containers in administrative groups; each container contains one or more routing groups. Then you assign servers to routing groups. When Exchange 2000 is in native mode, you can assign servers to routing groups regardless of which administrative group they belong to. Remember that assigning a server to a routing group in a different administrative group does not change the administrative group the server belongs to.

Create routing groups to match your network topology:

- If Exchange is in mixed mode, consider creating routing groups that mbdivide the routing structure for the Exchange 2000 servers that belong to large Exchange 5.5 sites.
- If Exchange is in native mode, consider creating an administrative group with a routing group container that holds all Exchange servers. This allows a single view of the routing configuration, regardless of the administrative groups individual servers belong to.

System Policies

System policies provide a method for administrators to apply settings to a group of servers simultaneously. You can create three types of system policies:

- Server
- Mailbox store
- Public folder store

After an attribute has been configured by a policy, it is no longer available in Exchange System Manager. You cannot change it except by modifying and reapplying the policy or deleting the policy object.

You can create system policies in one administrative group and apply them to servers in other groups. That way, one set of administrators can control particular settings through policies and delegate rights over other settings to administrators who handle other administrative groups.

Use system policies to ensure consistent application of settings to all servers. If you use a delegated administrative model, but need to apply control over corporate-wide standards, such as mailbox size or message retention times, consider using a centrally controlled administrative group that holds system policies.

Exchange Administration Delegation Wizard

Similar to Active Directory Delegation of Control Wizard, the Exchange Administration Delegation Wizard is a tool that assists in delegating responsibility for managing Exchange configuration objects in Active Directory. As described in Table 11.3, the tool allows you to assign one of three roles to groups or individual users. These roles are similar to the roles assigned in Exchange 5.5. Use the wizard to modify the organization or the individual administrative groups. Permissions are inherited from the top down, but can be overridden at lower levels if desired.

Table 11.3 Administration Delegation Wizard roles

Role	Right
Exchange View Only Administrator	View configuration of all Exchange objects.
Exchange Administrator	Modify configuration of all Exchange objects.
Exchange Full Administrator	Same as Exchange Administrator with the added ability to modify permissions on Exchange objects.

Create groups for the purpose of delegating administration by using Exchange Administration Delegation Wizard instead of specifying individual accounts. This makes it easy to change administrative rights by changing group membership rather than having to run the wizard again.

Additional Administrative Features

Whether you choose a centralized, delegated, or messaging group model, all administrators need to organize users into distribution lists and design public folder hierarchies. It may also be necessary to change the association of some users to a different administrative group. You can move mailboxes from one administrative group to another using Active Directory Users and Computers. This section discusses the tasks and the Exchange features that make it easy to manage recipients and servers.

Security and Distribution Groups

In Exchange 2000, mail-enabled groups held in Active Directory have replaced Exchange 5.5 distribution lists. There are two types of groups in Active Directory: security and distribution. Both types can be made mail-enabled and used as the equivalent of distribution lists. Security groups are used to control access to resources and can also be used as e-mail distribution lists. Distribution groups are used for e-mail distribution list purposes only. You can make a security group mail-enabled to inform members of changes to the resources that the group controls, and to simplify distribution list design by avoiding the creation of duplicate distribution groups. Distribution groups allow you to create distribution lists without the possibility of accidentally granting access to secure resources. Use mail-enabled security and distribution groups to create distribution lists for various groups within your company.

A suggested policy for mail-enabling groups is:

- Use security groups as the primary type, but make the groups mail-enabled only when appropriate.
- Use distribution groups for lists that include non-trusted recipients.
- Use universal groups when the ability to view membership is important. In Windows 2000 mixed mode, you need to use distribution groups.

You can further divide each type of distribution group into three categories, depending on whether you want the group to be accessible to specific domains or to all users. Exchange 2000 servers can route mail to group members regardless of the distribution group or category. The group type is more relevant when you consider how Outlook users will view groups and how you grant access to secure resources. For more information on recipients and groups, see the Windows 2000 and Exchange 2000 documentation.

Public Folders

Administration of public folders has changed from Exchange 5.5, including access control and public folder affinity.

Access Control

In Exchange 5.5 and earlier versions, access to public folders is controlled by granting permissions to mailboxes and users. Access to mailboxes is controlled separately, by granting permissions to Windows NT 4.0 user accounts. With Exchange 2000, the security model is unified. Access to public folders and all other aspects of Exchange is granted directly to user objects and group objects in Active Directory.

When Exchange 2000 coexists with Exchange 5.5, the group types used to control access to public folders are important. Active Directory Connector synchronizes Exchange 5.5 distribution lists into Active Directory as universal distribution groups. You cannot use these groups to control access to Exchange 2000 public folders. Exchange 2000 can convert universal distribution groups to universal security groups when folders are upgraded or replicated from Exchange 5.5 to Exchange 2000, or when the Outlook client is used to assign rights to an Exchange 2000 public folder.

Examine public folder permissions carefully when you migrate from Exchange 5.5 to Exchange 2000. Though this is not required, it simplifies the migration and testing process if particular public folders are housed exclusively on Exchange 5.5 or Exchange 2000. Permissions should be reassigned when you move folders from Exchange 5.5 to Exchange 2000.

Public Folder Affinity

Microsoft Outlook clients connect to Exchange servers to view the public folder hierarchy. When users attempt to open a folder that is not located on their server, the server provides Outlook with a sorted list of other servers that house a replica of the folder. Exchange 2000 uses routing groups and the cost on the connectors between them to determine the closest replicas. By default, public folder referrals can occur across connectors, but you can disable this feature.

A key difference from Exchange 5.5 is that in Exchange 2000, affinities are transitive. With Exchange 5.5, specifying affinity between sites A and B, and between B and C, does not automatically provide affinity between A and C. With Exchange 2000, if connectors exist that allow mail to be routed between all locations, all servers receive public folder referrals by default.

Moving Mailboxes

With Exchange 2000, you can move mailboxes by using Active Directory Users and Computers. You can move all mailboxes from one server in an administrative group to a server in another administrative group; in effect, you are moving the mailbox store of the server. Moving mailboxes is illustrated in Figure 11.1.

To move a mailbox

1. Open Active Directory Users and Computers.
2. Right-click the user object whose mailbox you want to move, and then click **Exchange Tasks.** You can select multiple users for batch moves.
3. In Exchange Task Wizard, click **Move Mailbox,** and then follow the directions.

Note This method does not move the public folder stores on a server.

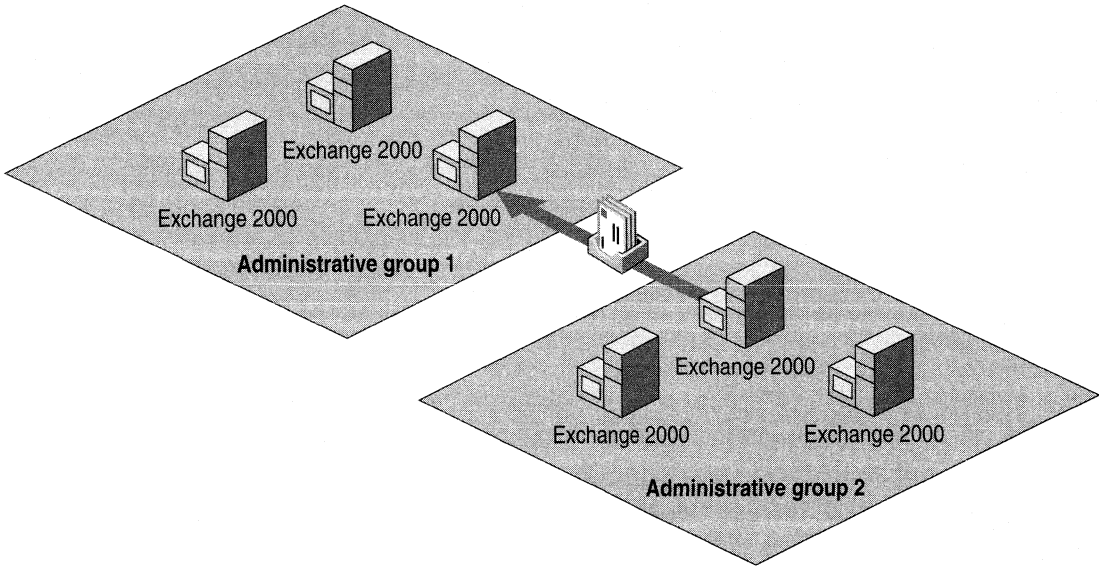


Figure 11.1 Moving mailboxes from one Exchange 2000 server to another

Consider the following three issues before you move mailboxes from one Exchange 2000 server to another:

- Additional hardware
- Network bandwidth
- Group policies

Additional Hardware

When you move mailboxes from one Exchange 2000 server to another, the target server should have similar or greater hardware capacity, such as dick space and memory, than the source server.

Network Bandwidth

A large amount of bandwidth is required to transmit tnailbox data between servers. The amount of bandwidth consumed when moving mailboxes is equal to the sum of the individual users' mailbox sizes—not the size of the database. This number can be difficult to calculate, but you can generate a rough number by using performance counters in Microsoft Windows 2000. In the System Monitor utility of the Performance console, examine the MSExchangeIS Mailbox Single Instance Ratio counter. Multiply the number provided by this counter by the database size to get a rough estimate of the required bandwidth.

Group Policy

When you move mailboxes from one Exchange 2000 server to another in a different administrative group, policies might vary between the two groups. You should consider the following issues when setting Group Policy:

- General policy
- Database policy
- Limits

General Policy

On the **General Policy** tab, you can set three items, any of which can cause problems during a mailbox move from one administrative group to another. Set the following options carefully:

- **Default public store** The default public folder store can vary from administrative group to administrative group. When you move mailboxes to a server in another administrative group, users could lose or gain access to certain public folders. It is important to address this when you move mailboxes.
- **Offline address book** The offline address book can be different among administrative groups. This can cause unexpected slowdowns for users. For example, if a user downloads an update to the address book after his mailboxes are moved to another administrative group, and the offline address book does not match his own, the entire address book is downloaded to his workstation. Over a low-bandwidth line, downloading the whole address book can be very time consuming. Although this is not likely to cause downtime for users, it can take considerable time to download the new information; this could generate a lot of calls to your Help desk.
- **Convert inbound Internet messages to fixed font** Messages from the Internet that were previously in Rich Text Format (RTF) or HTML can be converted to a fixed font after a mailbox is moved to a new administrative group. Conversely, messages from the Internet that were previously in a fixed font can be converted to RTF or HTML. You can control message conversion at the organization level by creating virtual Simple Mail Transfer Protocol (SMTP) domains in Internet message formats. If you want a server to deviate from an organization-wide policy on message conversions, you can also create Web Storage System policies for that server. Although font changes in messages might not adversely affect users, users might notice the difference.

Database Policy

The database policy can be different among administrative groups. The only issue you must consider here is the server maintenance schedule. Changes to the schedule could cause slower response times; however, they should not cause any downtime.

Limits

You set general limits in an administrative group policy. There are two different general settings in a limits policy:

Deletion settings The deletion settings policy can vary between administrative groups. The two administrative group-wide policies are:

- **Keep deleted items for X days.** If this setting is changed, users whose mailboxes have been moved could have items deleted from their Deleted Items folder on a different schedule than they originally set.
- **Don't permanently delete mailboxes and items until store has been backed up.** If this setting differs from the original administrative group, items might not be recoverable or the database size might increase beyond acceptable recoverable limits.

Storage limits When you move users from one server to a server that has a different storage limit, users might not able to send or receive messages because their mailboxes exceed the new storage limit.

This relatively simple item is easy to overlook when you move mailboxes between administrative groups. However, it can cause downtime for your users.

Mixed Mode Operation with Exchange 5.5

If Exchange 2000 operates in mixed mode with Exchange 5.5, there are some special considerations you should keep in mind. For example, it is recommended that all Exchange 5.5 servers continue to be managed from the existing Exchange 5.5 administrator program, and that all Exchange 2000 servers be managed from the Exchange 2000 System Manager.

Problems might occur if you attempt to manage Exchange 5.5 from Exchange 2000 System Manager or if you attempt to manage Exchange 2000 from the Exchange 5.5 administrator program. These issues are related to the fact that Exchange 5.5 does not recognize Active Directory, and if your system does not have a Site Replication Service, the Exchange 2000 server does not recognize the Exchange 5.5 directory. For example, if you try to access Exchange 5.5 mailboxes from Active Directory Users and Computers, the tool won't resolve the mailboxes the way that the Exchange 5.5 administrative tool does. Subsequently, you won't be able to move mailboxes that exist on Exchange 2000 servers by using the Exchange 5.5 administrative tool.

Administrative Groups

In mixed mode there is one administrative group in Exchange 2000 for each Exchange 5.5 site. Although it is possible to create more than one routing group within each administrative group, all routing group members must belong to the same administrative group as the routing group.

Active Directory Connector

The Active Directory Connector (ADC) is the tool that synchronizes the Exchange 5.5 directory with Windows 2000 Active Directory. For more information about ADC, see the Exchange 2000 Server online documentation.

Public Folders

Exchange 2000 supports item-level security in public folders, whereas Exchange 5.5 does not. When using this feature, make sure that public folders are not replicated to Exchange 5.5 servers. It is also important to make sure that Exchange 5.5 distribution lists are synchronized into native mode Active Directory domains so that Exchange 2000 can convert the lists to universal security groups and can grant access to public folders correctly.

Moving an Exchange 5.5 Server

Figure 11.2 illustrates the conditions required for moving an Exchange 5.5 server to a mixed-mode environment. To move an Exchange 5.5 server to a mixed-mode environment, the Exchange 5.5 server being moved must be running on a site containing only Exchange 5.5 servers (organization 1). Also, there must be at least one Exchange 5.5 server in the destination site (organization 2), which is also an Exchange 2000 administrative group.

Note You cannot install an Exchange 5.5 server in a native mode Exchange 2000 organization.

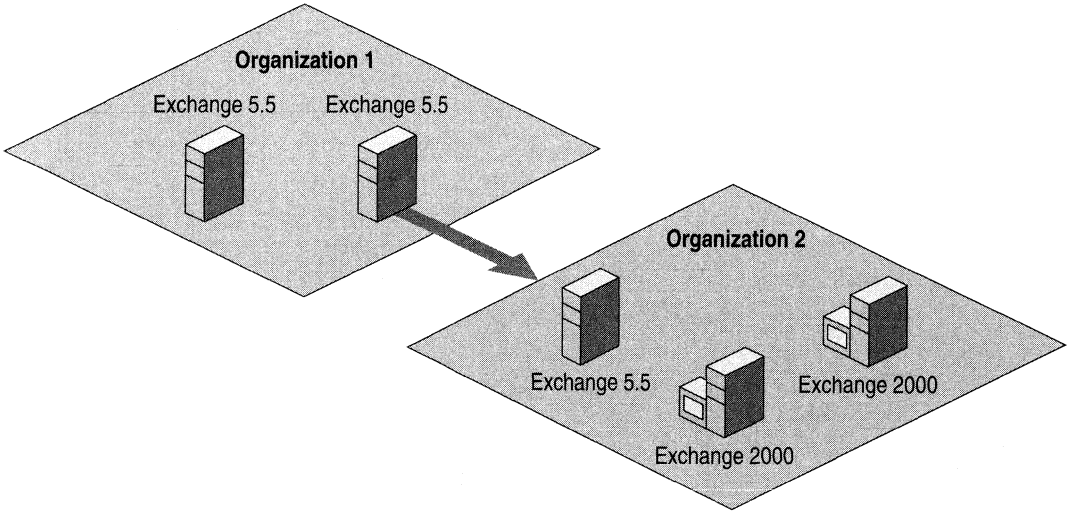


Figure 11.2 Moving an Exchange 5.5 server to a mixed-mode organization

To move an Exchange 5.5 server to a nixed-mode environment, you use the Exchange 5.5 Move Server Wizard. This wizard allows you to move the server from one Exchange 5.5 site to another. The same issues—such as connectors and public folders—that arise in an Exchange 5.5 only environment also arise in a mixed-mode environment. In addition, you must also consider Active Directory directory service connection agreements before you move a server.

If the server that you are moving has an Active Directory connection agreement, you must move this connection agreement to another server on the original site before you move the server. If you don't, and this is the only connection agreement on this site, you will no longer be able to administer the site from Active Directory Also, in some circumstances you can orphan all recipient objects in Active Directory for this site. If this happens, you can remedy the situation by re-creating a connection agreement in the original site. For additional information, see Microsoft Product Support Services Knowledge Base article Q196413, "XADM: Remove Connectors Before Running Move Server Wizard."

Managing and Monitoring Exchange Server Performance

Managing and monitoring an Exchange organization is important for three reasons:

- **Capacity planning** Normal operation statistics about the system are valuable for predicting how the system will function as it grows in size. This allows accurate planning for modifications and enhancements.
- **Problem notification** Monitoring tools can alert you of problems early on. You can then resolve the problems quickly, before they have an impact on users.
- **Troubleshooting** When a problem arises, management and monitoring tools provide vital insight into the system and allow for efficient isolation of problems. You can identify the root cause and correct the problem.

When you plan a new installation or upgrade an existing one, it is important not to concentrate solely on Exchange 2000 architecture design issues. Management and monitoring decisions should not be neglected or left for later. Rather, management and monitoring should be considered an integral part of the initial design effort. It is much easier to build a good set of tools into a new installation than to retrofit an installation later.

Exchange 2000 Tools

Microsoft Exchange provides both *reactive* tools, which help you manage system problems after they have been reported, and *proactive* tools, which search continuously for potential problems and take automatic corrective action when they are found. Exchange 2000 includes three key tools for monitoring and troubleshooting:

- Exchange Monitoring and Status Tool
- Queue Viewer
- Message Tracking Center

Exchange Monitoring and Status Tool

The Exchange Monitoring and Status Tool is located under **Tools** in Exchange System Manager in MMC. It replaces the Server Monitor and Link Monitor tools in previous versions of Exchange. This new tool has the ability to provide notifications based on the continued growth of queues rather than a static threshold.

All servers and connectors in the organization are listed in Exchange System Manager, in the Status container; a filtering option lets you limit which servers and connectors are displayed. Each connector is automatically monitored and is shown to be available or unavailable. For each server, one or more resources can be monitored. By default a resource called Default Exchange Services is monitored. It includes the following services:

- Microsoft Exchange Information Store Service
- Microsoft Exchange MTA Stacks
- Microsoft Exchange Routing Engine
- Microsoft Exchange System Attendant
- Simple Mail Transport Protocol (SMTP)
- World Wide Web Publishing Service

You can add other services to the Default Exchange Services resource or to the Windows 2000 service resource. Table 11.4 gives a complete list of resources that can be monitored.

Table 11.4 Resources to monitor

Resource	Purpose
Default Exchange services	Monitors specified services. Can be changed to warning or critical state if a service is not running.
Available virtual memory	Monitors available virtual memory with control over the thresholds for warning or critical state and how long the threshold must be exceeded to change the state.
CPU utilization	Monitors available CPU utilization with control over the thresholds for warning or critical state and how long the threshold must be exceeded to change the state.
Free disk space	Monitors available disk space on a specified drive with warning and critical thresholds.
SMTP queue growth	Monitors growth in SMTP queues with thresholds in minutes for warning or critical state when queues grow continually.
Windows 2000 service	Monitors specified services. Can optionally change to warning or critical state if a service is not running.
X.400 queue growth	Monitors growth in X.400 queues with thresholds in minutes for warning or critical state when queues grow continually.

When any of the monitored resources for a server changes state, the monitor will display a warning or critical state. Exchange administrators can create e-mail, page, or script notification entries in the Notifications container that are triggered when the status of servers or connectors changes. Available parameters for configuring notifications are shown in Table 11.5.

Table 11.5 Parameters for configuring notifications

Parameter	Specific Item to Monitor
Monitoring server	Indicates the server that performs the monitoring.
Servers and connectors to monitor	Indicates whether this server, all servers or connectors, any server or connector in the routing group, or a customized list of servers or connectors are to be monitored.
State	Controls whether notification occurs on status change to warning or critical.

When a notification is triggered, three types of notification can be sent to the administrator, depending on the parameters or details that are available. Table 11.6 lists the notification types and the details needed to send the specific notification types.

Table 11.6 Information required for notifications

Notification Type	Parameters
E-mail	Recipients, sending server, subject line, and message body.
Page	Pager number and touch-tone message to send.
Script	Path to executable and command line switches to use.

Use the Exchange Monitoring and Status tool to establish comprehensive monitoring of all Exchange 2000 servers and connectors. Perform baseline analysis during the rollout of Exchange and turn on notifications gradually to ensure that triggers are set to appropriate levels and to avoid having the system generate alerts more often then necessary.

Queue Viewer

Exchange 2000 System Manager includes a new Queue Viewer tool that you can use to view the state of all queues. Messages can be frozen in the queue, returned to the sender, or deleted.

Use this tool for problem isolation and troubleshooting after the Exchange Monitoring and Status tool highlights a problem.

Message Tracking Center

Although Exchange 2000 has a new message tracking log format, the features of the Message Tracking Center are similar to previous versions. You can use the tool to track messages from sender to recipient through the Exchange system.

Enable message tracking for all Exchange components. Whereas message tracking can have a slight impact on performance, the disadvantages are minor considering the power of today's server hardware and when compared to the benefit of investigating user complaints regarding missing or delayed e-mail. Without the information provided by message-tracking logs, support personnel have little choice but to trust the reports of end-users. It is critical to have access to reliable data on system performance, particularly when the system architecture is going through a change.

Windows 2000 Tools

Windows 2000 includes many tools useful for managing the system. Two of these tools, System Monitor and Event Viewer, are described in this section.

Performance Monitoring

Windows 2000 and Exchange 2000 contain numerous performance counters. The Performance Logs and Alerts snap-ins for MMC allow you to record system activity over a period of time for analysis later. In addition, you can configure alerts so activity that crosses defined thresholds generates events, such as adding an entry to the application event log. sending a message to a Windows 2000 console, starting a performance data log, or starting a program.

System Monitor, available by choosing Performance from the Administrative Tools folder, can chart activity in real time. or display information contained in a log file.

These tools are useful in the following areas of Exchange administration:

- Real-time alerts on the performance of selected counters, such as the MTA Work Queue Length, which is a good indication of the health of the Exchange message transfer agent (MTA). The Exchange 2000 monitoring tool also covers several important real-time measurements.
- Real-time analysis of a server that is experiencing problems, so that the failing component can be isolated.
- Capacity planning and trend analysis, which can be performed using data logged continually to disk.

Capture important statistics from all servers for use in capacity planning and to determine if problems are caused by changes over time.

System Monitor should also be used when a server is exhibiting problems. A quick analysis of processor activity and queues can be useful in isolating problems with specific components. It is important to be patient when servers arc recovering from outages, because services may be working hard, for example to rebuild queues, without producing noticeable results. Monitoring performance activity is key to determining whether or not the services are working.

Event Viewer

Like all other applications running on Windows 2000, Exchange 2000 logs all error messages and warnings to the Windows 2000 application event log. Messages specific to Windows 2000 operation are entered in the Windows 2000 application event log. You can view and filter these errors through Event Viewer in MMC. A central log of this sort can help you diagnose problems, particularly when all system errors are logged to the same place.

Exchange 2000 logs critical events to the Windows 2000 application event log. If necessary, you can configure each Exchange component to generate high levels of logging. Under normal conditions these logging levels should be set to their minimum level (None) to avoid filling the event logs too rapidly. When you suspect that a component has a problem. you can increase the level of logging.

You use Event Viewer to view the event logs manually. The *Windows* 2000 Server ***Resource Kit*** companion CD includes utilities that you can use to export the content of event logs to files for subsequent analysis.

System administrators should establish a policy of manually scanning Exchange Server event logs on a regular basis. Being familiar with normal event logs gives you valuable insight when you examine the event log of a server that is experiencing problems. In addition, you can monitor the system for problems by using automated procedures.

CHAPTER 12

Server Design for Backup and Restore

Peter Nilsson, Principle Consultant, Microsoft

Backup and restore is an important topic for Microsoft Exchange 2000 Server. Fundamental design changes in Microsoft Web Storage System make it possible to distribute data across multiple databases (mailbox stores and public folder stores) and storage groups. This distribution of data can decrease your backup and restore times provided that your hardware is properly configured.

Important issues and common scenarios appear at the end of this chapter to provide ideas for how you can design a backup and restore strategy for your environment.

Because substantial documentation exists for backup and restore for Microsoft Exchange version 5.5, this discussion focuses on issues that are unique in Exchange 2000. For additional technical information about backup and restore, see "Backup and Restore" in this book.

In This Chapter

Databases and Storage Groups

Design Issues

Sample Designs

Databases and Storage Groups

It might be easiest to understand backup and restore in Exchange 2000 by looking at the differences between Exchange 5.5 and Exchange 2000. For information about backup and restore in Exchange 5.5, see the Exchange 5.5 documentation and related sources on the Exchange Web site at http://www.microsoft.com/exchange.

The first and most obvious difference is that the architecture of the Web Storage System and Extensible Storage Engine (ESE) has changed:

- Instead of a fixed database layout, Exchange 2000 has storage groups. A storage group corresponds to an instance of ESE (with it's own sequence of transaction log files). Exchange 2000 supports four storage groups per sewer.

  Note In this discussion, the generic word *database* refers to either a mailbox store or a public folder store in a storage group.

- Each storage group can support up to five databases, and each database can contain either mailboxes or public folders. The transactions for all databases in a storage group are contained in the single set of log files.

Although the backup application programming interface (API) incorporates a number of changes, online backup still looks and acts very much the same as in Exchange 5.5. On a server with a single mailbox store or public folder store, the differences are minimal. Most changes pertain to backing up multiple storage groups and databases.

You can restore a single database in a storage group running multiple databases without taking the others offline. You can run parallel backups and restores to support large configurations. These features are possible because:

- Backup runs on a storage group. This is because a storage group corresponds to an instance of ESE and an instance of the backup API.
- Backup runs sequentially against the databases in a storage group. It is not necessary to back up all the databases in a storage group as part of the same job; old transaction logs are not purged until all databases have been backed up. After a full backup, two things are deleted: the transaction log files, and the transactions that have been committed to the databases and that have been backed up. Incremental backups will back up and delete transaction logs before the checkpoint file.
- Storage groups can be backed up in parallel. Each storage group is an instance of ESE, and these run independently of each other, at least as far as backup is concerned.
- A database can be restored without affecting databases running in the same storage group. To do so, initialize a reserved instance of ESE to handle the restore (the Web Storage System is able to support more storage groups than can be created on an Exchange server; the additional capacity allows for this reserved instance of ESE). You can restore the database using this temporary instance of ESE, then dismount the database and mount it in the correct storage group.
- Databases can be restored in parallel.

These features of Exchange 2000 make backup and restore designs and associated procedures more complex than in Exchange 5.5. There are some detail changes around restore and recovery that are necessary to support this level of potential complexity.

- A restore-in-progress key is no longer used during restore. Individual data structures are created for each database being restored.
- It is no longer possible to allow recovery without solving corruption problems, for example, to play the wrong transaction log files into a database, or to trick a database into starting when the necessary components are not present on disk.

Backup

Backup works for Exchange 2000 in basically the same way as Exchange 5.5. However, there are a few differences:

- Each database consists of two files: the .edb file and .stm file. They are backed up together. The backup process continues sequentially until all of the databases in the storage group that have been selected for the current backup are copied to the backup device.
- The transaction log files and patch files have checksums that are validated during the backup process.
- The transaction log files are not truncated until all databases in the storage group have been backed up. After a full backup, two things are deleted: the transaction log files and the transactions that have been committed to the databases and that have been backed up. Incremental backups back up and delete transaction log files that precede the checkpoint file.
- A database must be online to be backed up. If a database has been dismounted it cannot be backed up, and the transaction log sequence will not be truncated.
- Conduct a full backup after switching from circular logging to non-circular logging. During circular logging, information in the .stm file is not recorded in the log files. When you change to non-circular logging, transaction log files still exist that do not have .stm file data; these transaction logs must not be replayed.

Restore

Restore in Exchange 2000 has changed more than backup. Before you can attempt to restore a database in Exchange 2000, the following must be true:

- The relevant service and the Web Storage System must be running.
- The database to be restored must be dismounted.

The significant differences are as follows:

- It is possible to restore multiple databases from the same storage group as part of a single restore job. In this case, the restore process restores all of the databases to disk before continuing.
- The transaction log files in the backup set and the patch files are restored to the temporary disk location specified by the user. The information about the restore previously written to the restore-in-progress key is written to a file called Restore.env.
- If multiple datasets are being restored (for example, for differential or incremental backups), the dataset containing the full backup must be restored last. When the last dataset is being restored, you must select **Last Backup Set.**
- After all files are restored, recovery begins. The Restore.env file is used to find the end and beginning transaction log numbers and the relevant transactions are replayed into the database. After the end log is replayed, recovery goes to the transaction log files of the target storage group and continues to play through additional log files until the end of the sequence is reached.
- After restore finishes, the database is dismounted from the temporary instance of ESE and the files in the temporary work area are deleted. If you selected **Mount Database after restore,** the database is automatically mounted in the target storage group.

Because parallel restores are possible, the restore process relies on the user to provide a path to temporary disk space that will be used during the restore. Separate restore processes running at the same time must use different disk locations. The temporary disk space required is about 10 megabytes (MB) more than the size of the transaction log files and patch files that are being restored.

After the relevant files are restored to disk, the backup process will have to replay logs to process log and patch files and make the database consistent. An instance of ESE is required to perform the recovery, and this is where the reserved instances of ESE are used. ESE itself can support 16 instances, whereas the Web Storage System does not, so there are enough instances of ESE to run recoveries in parallel.

Parallel Operations

Parallel backups and restores put far more stress on input/output (I/O) subsystems than single database backups and restores. Particular attention must be paid to aggregate I/O bandwidth over the entire data path between database disks and backup devices.

Exchange 2000 backup and restore is faster than previous versions of Exchange. Rates of up to 70 gigabytes (GB) per hour on backup and 40 GB per hour on restore are possible. However, you must carefully design the data paths in order to support several concurrent operations at these rates.

Server Recovery

Thus far, the discussion has covered the need to restore one or a few of the databases that might be running on a server. I-Iowever, your backup design must be able to cope with a situation where all databases have to be restored. This scenario is traditionally part of a major disaster plan that accounts for an entire server or computer room that has been seriously damaged. This type of scenario might also occur if a serious accident causes major damage to an entire disk subsystem. With the growing trend towards Fibre Channel-attached external storage, problems can affect multiple servers.

The key to a complete and quick recovery is that, as your servers grow in capacity and as you create and support more databases per server, you need to maintain sufficient capacity to handle the maximum load a total restore might place on your server. You can either expand the server's capacity for data throughput, or take the risk that your system might not meet restore requirements under all circumstances.

Recommendations

As with earlier versions of Exchange, the primary recommendation with Exchange 2000 databases is to perform full online backups.

- Only a full online backup verifies the checksum on every page in the database.
- Restore from a full backup represents the fastest way to recover with full confidence from a corrupt database.

Ideally, a full online backup should be run every night on every Exchange database.

With Exchange 2000, it is a good idea to back up all of the databases in a storage group at the same time. This will ensure the most regular and predictable backup of transaction log files, and also ensures that all of the data that might be required to restore a storage group will be in one place.

Design Issues

This section contains information about backing up large quantities of data as efficiently as possible. You can back up to disks or to digital tape, depending on the volume of your backups and acceptable restore times.

Backup to Disk

The Microsoft Windows 2000 Backup program can backup Exchange databases to disk. Although this is useful especially for testing, certain limitations make it impractical to base regular backups on this approach.

Backing up to disk is not advantageous. It might seem as if a backup to disk would enable a quick restore, because the files you want to restore are already on disk. However, all of the files associated with an online backup, such as databases, log files, and patch files, are stored in a single backup file; they must be restored to disk from this file in the same way as you would restore from any other backup media. Disk backups can be slower because disks are not efficient streaming devices. You can stream to a single, fast tape device, such as a digital linear tape, far more quickly than you can stream to a single disk. Multiple drives and caching disk controllers can improve performance, but a Redundant Array of Inexpensive Disks (RAID) decreases performance, and it is difficult to build a disk array that can write as quickly as a single tape, let alone an array of tape devices.

Backup and Restore Throughput

High throughput for backup and restore depends on your hardware configuration, and thus remains the same for Exchange 2000 as they were for Exchange 5.5.

This section does not discuss detailed hardware configuration issues; rather, it provides a high-level discussion of some general hardware issues as they relate to throughput.

- Tapes are faster streaming devices than disks, but they slow dramatically if the streaming speed cannot be maintained. Therefore, you need to initiate backups from multiple disks simultaneously to provide sufficient flow of data to keep tapes at streaming speeds. You also need to ensure that adequate bandwidth is available between the disks and tapes.
- Hardware compression on the tape device boosts backup speeds. For example, the maximum speed you can write to a digital linear tape 35/70 GB equates to around 20 GB per hour. A typical Exchange database can usually be backed up at 25 to 30 GB per hour, although the physical speed of writing to the tape remains the came (20 GB per hour). The device is able to accept data at a faster rate and compress it into a 20–GB-per-hour stream to the tape.
- If one tape isn't fast enough, it is often possible to write to multiple devices in parallel. Previously this was achieved by using only the features of your backup software and hardware because an instance of ESE provided a single backup stream. With Exchange 2000, it is possible to back up multiple storage groups in parallel; however, to make a single storage group backup faster, you still need to look to a backup software or hardware solution.
- Restores are almost always slower than backups. This is usually caused by a combination of disks being slower streaming devices than tapes, the transfer rate for the RAID level used for your database partition, and features of the disk controller. Typically this means that you need to restore to multiple disk drives in parallel. The more tape devices you run in parallel, the more drives you need.

RAID-5 in particular presents some well-known performance challenges, because parity must be calculated and written with each write. Some high-end RATD controllers can recognize that they are streaming large blocks of data to disk; if complete disk blocks are being written with each write operation, there is no need to read the parity block before each write. This optimization can provide a significant performance improvement.

Generally speaking, high-end controllers will perform resource-intensive operations such as bulk restores better than cheaper controllers, but one feature is essential: a well protected write-back cache. Without a write cache to streamline the flow of data to disks, restore rates can slow to half the normal speed.

The recommended backup strategy is to perform a full online backup of all databases every night.

You need to ensure that the backups do not overlap with system maintenance, or online defragmentation may never run. If backups for a storage group begin while a database in the storage group is being defragmented, the defragment process ceases.

Larger server configurations require more complex backup designs as you try to funnel more data over limited bandwidth within a limited interval.

Sample Designs

This section applies some of the principles discussed thus far to typical deployment scenarios. One scenario is presented for small organizations or branch offices, a second for a departmental server, and a third for centralized server or data-center server.

Small Organization or Branch Office Servers

First consider the logical layout of a typical small organization or branch-office server. Simplicity is the key to a low-cost design. Figure 12.1 shows a simple Exchange 2000 server configuration.

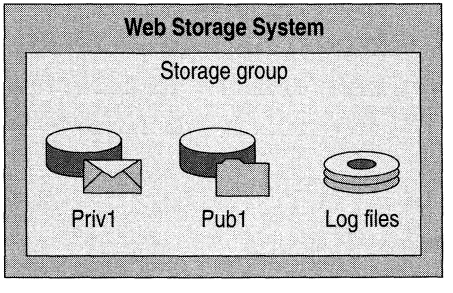


Figure 12.1 Simple Web Storage System design

The layout resembles an Exchange 5.5 server, and the physical design of the server would include the following:

- Mirrored system disk, which contains the operating system, application binary files, paging file, and so on.
- Mirrored transaction log disk.
- RAID-5 database partitions. For the purpose of this example, assume that there are six 9–GB drives that provide 45 GB of usable space.

  Note Six 9–GB drives have been chosen instead of three 18–GB drives. Fewer drives provide less usable space (only 36 GB), and are probably not able to produce the transfer rates needed to keep a digital linear tape streaming, resulting in significantly slower backup.

Because tapes have an uncompressed capacity of *35* GB, and the database partition should never be filled to more than 75 to 80 percent (according to operational best practices derived from Exchange 5.5), you can be reasonably safe in assuming that, for this example, you will always be able to back up the database partition onto a single digital linear tape 35/70 GB. This means you don't need to change tapes during the nightly backup, and you can deploy the simplest and cheapest hardware solution: a single small computer system interface (SCSI)-connected digital linear tape drive.

For this example, the backup should finish in just over an hour

Note that this example only provides a baseline for developing further designs. You still need to resolve who will change tapes during the day and how the system will be recovered.

Large Departmental Server

A server that supports approximately 2,000 mailboxes with a 50 MB per user quota will require about 100 GB of storage per server (calculated by multiplying the number of users with their maximum storage limit). Additional storage might be required for public folder storage. If, for example, the backup and restore design must provide for a database restore within two hours, and you want to localize the effect of errors and failures as much as possible, a potential configuration follows:

- By using a RAID-5 disk layout, you can expect restore rates of about 20 GB per hour. To meet the time constraint, your maximum database size is limited to 40 GB; to meet your capacity requirement, you'll need three mailbox databases (100 GB divided by 40 GB). By dividing mailboxes evenly between the three databases, you will achieve a maximum database size of 34 GB, which will accommodate the two-hour restore goal.

  An additional database can be added for the public folder store. This assumes that 34 GB of public folder storage will suffice.

- To accommodate four 34-GB databases, you need to determine how many storage groups you want. Four storage groups, with one database per storage group, will make the databases as independent of each other as possible, and maximize your options when it comes to parallel backup and restore operations.

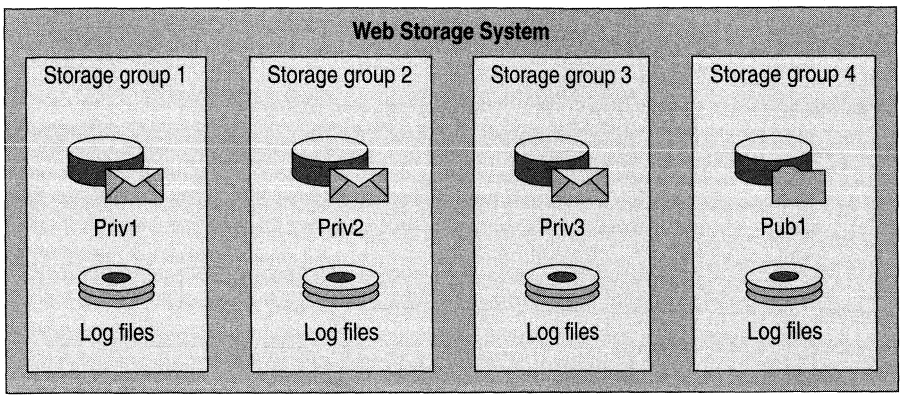


Figure 12.2 Storage group and database design

To provide good overall performance you should ensure that each database is spread over a reasonable number of drives, and also that each instance of ESE has a dedicated transaction log drive.

The overall disk layout is as follows:

- Mirrored system disk, which contains the operating system, application binary files, paging file, and so on.
- Mirrored transaction log disk for each ESE instance (8 disks total).
- Four RAID-5 database partitions, each made up of six 9-GB drives, providing a total of 45 GB usable space per partition.

In total you will need 34 disks, which you can support using three disk array controllers.

When it comes to designing the backup solution, you have a number of choices:

- You can run several simultaneous across-the-network backups to a robotic tape library.
- You can run several simultaneous local backups to SCSI-connected tape drives.
- You can run a single sequential backup to an array of SCSI-connected drives.
- You can run a series of separate sequential backups to a local SCSI-connected tape library.
- You can run a single sequential backup over meshed Fibre to a robotic tape library.

A large dedicated robotic tape library would be the fastest option if you have enough servers to justify the investment. The advantages and disadvantages for the three remaining choices for local backup are presented in Table 12.1. Because you calculate the total backup time to be four and one-half hours (based on backing up 136 GB at 30 GB per hour), you can factor backup speed out of the evaluation; sequential backups are necessary.

Table 12.1 Backup type advantages and disadvantages

Backup Type	Advantages	Disadvantages
Simultaneous backups to local tape drives	• Potentially the shortest backup time. • Each database is backed up on a separate tape. This should make restores and archive management easier. • Possible to run parallel restores; everything is in place to enable this.	• No resilience in backup media; exposure to tape failure. • Potential to swamp internal busses. • Manual tape changing required (during the day, not during the backup).
Sequential backup to a local array of tape drives	• Short backup time. • If using RAID-5 on the tape drives, you have resilience in the backup media.	• Only one restore at a time is possible (single backup set). • All databases are on one set of tapes; they can't be separated. • Manual tape changing required (during the day, not during the backup).
Separate sequential backups to a local tape library	• No daily operator intervention required. • Backups on separate tapes; easier management and storage.	• Longest backup time. • No redundancy in hardware media (unless RAID 1 can be used). • Only one backup at a time is possible (unless library has more than one tape drive).

Two of these issues are more crucial than the others:

- Redundancy in the backup media is important. Full backups purge old transaction log files without backing them up. so you can only replay log files after restoring the most recent backup. If you suffer a media failure on your most recent backup, you can restore an older backup but you can't replay its transaction logs and are likely to lose data.
- Unless you have Exchange server support at the location where the servers are housed, you must change tapes manually. In these situations, automatic tape libraries are well worth the financial investment.

Unfortunately, these two factors are not always compatible. For the previous example, the recommendation would be to go with media redundancy by installing an array of four digital linear tape 35/70 GBs and configuring them as a RAID-5 array in the backup software. In this case, the ability to perform parallel restores is sacrificed; clearly this represents a cautious approach that emphasizes the safety of backed up data. More aggressive approaches that enable parallel restores are equally feasible and may be more appropriate in different circumstances.

Centralized Server or Data Center Installation

To back up and restore databases on a 10,000-user server, consider the following configuration. Assuming the server supports 10,000 mailboxes with a 50 MB per user quota, the total required space on the server is 500 GB. For this example, assume that public folders are not part of the equation; this organization will employ separate, dedicated public folder servers.

If the backup and restore design must provide for a database restore within two hours and you want to localize the effect of errors and failures as much as possible, it would be best to keep the maximum database size to 50 GB, which equates to ten databases. You can set up nine databases (across three storage groups), and still assume a 50 GB maximum database on the grounds that theoretical size limits are rarely attained.

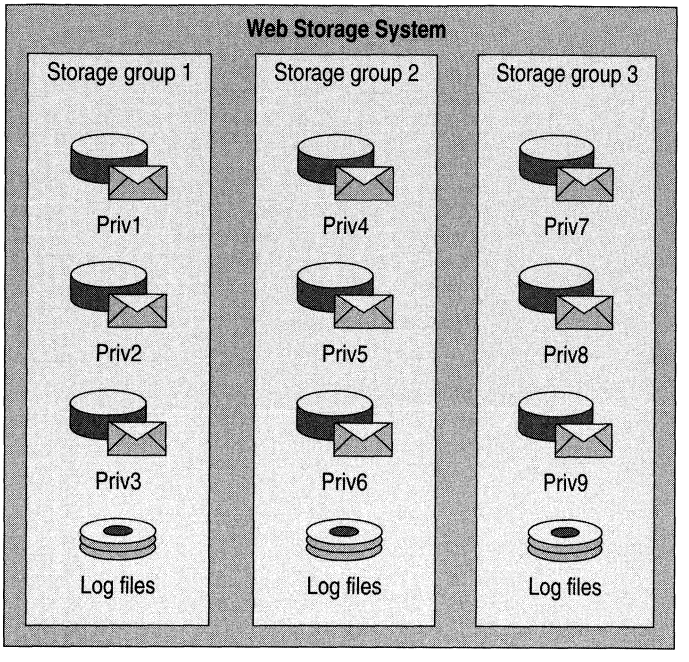


Figure 12.3 Storage group and database design for data center installation

Note The configuration here has not been validated. It is based on extrapolation from what is known about Exchange 5.5.

You must decide whether to put each database or each storage group into a separate partition. Each storage group will support 3,333 users, which is a significant load. You will want at least 12 to 15 spindles, preferable in a RAID 0+1 configuration, to provide sufficient disk throughput capability.

With a single partition for each storage group, you can choose between twenty 18–GB disks in a RAID 0+1 (180 GB usable) or twenty 9–GB disks in a RAID-5 (171 GB usable) arrangement.

With a single partition for each database, you would probably need three arrays of eight 9–GB disks in RAID-5 configuration (63 GB usable each), for a total requirement of 24 disks per storage group.

You will minimize the load on the disk subsystems by using a storage group configuration and putting each storage group on one RAID 0+1 partition, making this the better configuration in most cases. As such, you will need 60 18–GB disks for database partitions. This configuration leads to a specialized, high-capacity external disk storage system. With this type of solution you will also benefit from high-end array controllers, large caches, and high bandwidth that uses a switched Fibre Channel solution.

Each storage group also needs space for transaction log files. There are two main requirements for transaction log files:

- Physical separation from the database files
- Fast writes

It is recommended that you use four 18–GB disks in a RAID 0+1 array for all the transaction log files, instead of dedicated drives for each storage group. By using dedicated drives, you can ensure that both of the requirements are met. A high-end caching array controller will provide fast writes.

The backup strategy for this configuration is challenging. if you nave three disk partitions each with potentially 150 GB of data to back up, the total backup time at 30 GB per hour would be 15 hours. You could devise a complicated scheme of differential backups so that you don't back up every storage group every night, but you don't want to compromise the basic principles of your backup strategy in this way.

You need faster backup; for example, backing up to a Fibre Channel attached high-capacity tape library that writes to multiple tapes in parallel. With this configuration, your backup will run at 50 GB per hour and you will be able to back up a storage group to three separate tapes in three hours.

You can back up storage groups in parallel, or you can back up storage groups in sequence to reduce your bandwidth.

As discussed earlier, make sure that backups do not prevent defragmenting from running during system maintenance. You need to coordinate your system maintenance to make sure that all databases are defragmented.

Clearly such an example calls for a full-scale test, but this basic design offers a sound basis for a backup and restore scheme for a very large server.

CHAPTER 13

Virus Protection

William Riedell, Senior Consultant, Microsoft
Ramon Infante, Senior Consultant, Microsoft

During 1999, viruses transported by messaging systems cost businesses over 7.6 billion dollars. A recent industry survey concluded that computer infections have steadily increased over the past four years, raising the average infection rate to approximately 87 out of 1,000 computers. In addition, reports assert that between 5 and 40 percent of a given company's e-mail traffic is attributable to spam—unsolicited messages from retailers or other businesses.

This chapter provides best practices for using virus protection on servers running Microsoft Exchange 2000 Server. It is important to consider which protection methods are safe for maintaining system and information integrity but that still provide the high levels of protection needed to safeguard against attacks by computer viruses.

This chapter also describes the implementation of the virus scanning API in Exchange 2000, originally released in Microsoft Exchange Server 5.5 Service Pack 3.

In This Chapter

Virus Protection Overview

Solutions Planning

Virus Scanning Support in Exchange

Virus Scanning Support in Outlook

Virus Solution Vendors

Alert Monitoring

Virus Protection Overview

Over the past ten years, virus protection has primarily focused on solutions for the desktop. Companies and individuals have frequently updated signature files that scan for existing viruses on a local desktop. While this is fairly effective with viruses that are transferred via floppy disk, it is no longer an effective strategy. Today, viruses that spread through local area networks, Web sites, and e-mail can infect hundreds of thousands of users in a few hours.

It is important to protect your corporate environment at as many entry points as possible. Figure 13.1 illustrates a possible network environment. Remember that there are many potential virus entry points, including:

- Internet access
- Remote users and sewers
- LAN users who install software from diskettes
- Messaging traffic

These all offer a means of entry for viruses into the corporate environment; you should examine them carefully.

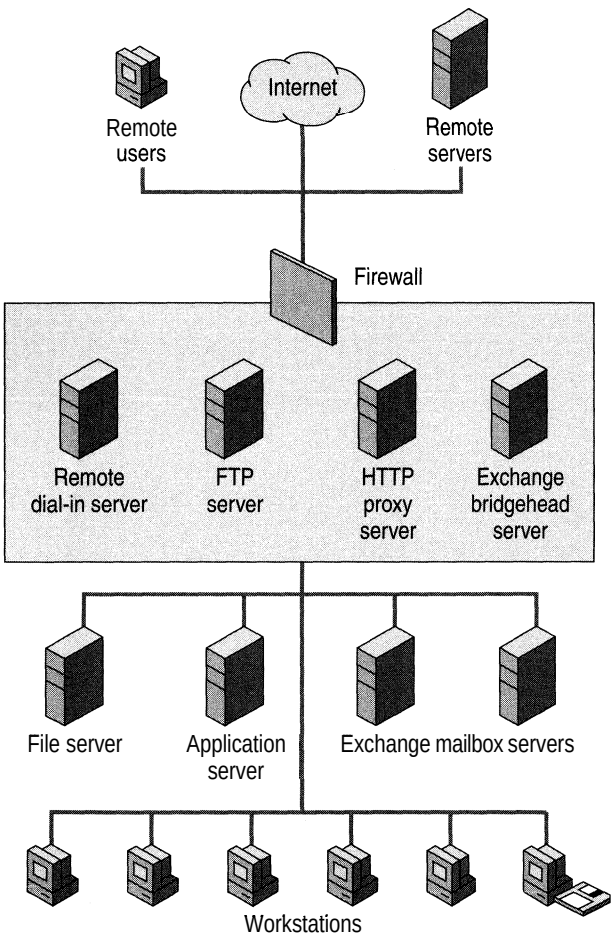


Figure 13.1 Virus points of entry

The growth of e-mail as a critical form of communication has increased the importance of developing and deploying a comprehensive virus-protection solution.

Today's antivirus measures require the ability to identify and stop threats before they reach desktop computers. An effective solution must include several components: a multi-tiered approach, a threat mitigation plan, a disaster recovery plan, and communication networks. You must also take into account the processes by which teams are established to help with antivirus protection. For example, some Microsoft customers have established internal virus control centers; these centers are responsible for the preparation and dissemination of all virus-related information, including sending alerts tc various departments and service centers, and setting standards for product and pattern updates.

Virus Scanning Concepts

Virus scanning software packages typically have two components: a search engine and a signature file. The search engine component searches memory and file storage, looking for viruses. The signature file provides descriptions of the viruses or virus patterns to the search engine.

Rather than requiring an exact match, the typical method of detection is to use a small piece of the virus code as an identifier. These short templates, called ***virus*** signatures. are easy to use and reveal nothing useful to virus authors. Using short signatures is advantageous because they remain useful even if other parts of the virus change.

A signature must satisfy two somewhat contradictory goals: it must capture a broad variety of conceivable mutations for a particular virus, and at the same time it must keep the false-positive probability as low as possible.

The virus protection package uses its search engine to look for viruses that match the virus signature file pattern. When a virus is found, the package takes action against the virus. The drawback to this method is that viruses tend to be created or modified at a much faster rate than signature files are updated. It is currently estimated that ten new viruses are discovered each day. This frequency makes it very difficult to deploy new signature files fast enough to large organizations where there may be hundreds of thousands of desktops.

The Exchange virus scanning API increases the chance of preventing virus attacks by scanning incoming messages using recent virus signatures before users receive messages at the desktop.

Definitions

- **Virus** A piece of code that replicates by attaching itself to other programs or files. When these files run, the virus is invoked and begins replicating itself.
- **Virus signature file** A file containing virus signatures that are compared with saved or incoming files to determine if the files are infected with a virus. The vendor of the antivirus software updates the signatures frequently and makes them available to customers over the Internet.
- **Virus signature** A binary pattern from the computer code for a particular virus. Virus scanning programs compare these patterns with the files on the hard disk, on removable media (including the boot sector5 of the disks), and in RAM.
- **Boot sector virus** A virus that is placed in the sector of a hard drive that is used to start the system. Once the computer is turned on or restarted, the virus is autoinatically executed.
- **File virus** A virus that attaches to an executable file. When the infected file is executed, the virus executes.
- **Macro virus** A macro or script that attaches to a file or template. Once the file is loaded, the macro or script instructions are executed.
- **Trojan horse** A piece of code embedded in a useful program for malicious purposes. A Trojan horse differs from a virus in that it does not try to replicate itself to other programs.
- **Worm** A program that replicates by running copies of itself across a network. A virus can exhibit both virus and worm characteristics, as in the recent case of the Explore.zip virus.
- **Spam** Unsolicited "junk" e-mail.
- **Vandal** An auto-executing virus written by using such technologies as Microsoft ActiveX controls, Java applets, auto-executable plug-ins, and Dynamic Hypertext Markup Language (DHTML).
- **Multipartite virus** A combination of the viruses listed above.

Solutions Planning

There is no one solution that allows an organization with distributed system architecture to combat viruses at a single location. The defenses must be implemented at numerous levels—from the routers and firewalls that provide basic connectivity to the desktop where end users perform their work. You do not need to purchase solutions from a single vendor for all three tiers.

Figure 13.2 illustrates how virus protection fits into the three tiers; this is described in detail in the following sections.

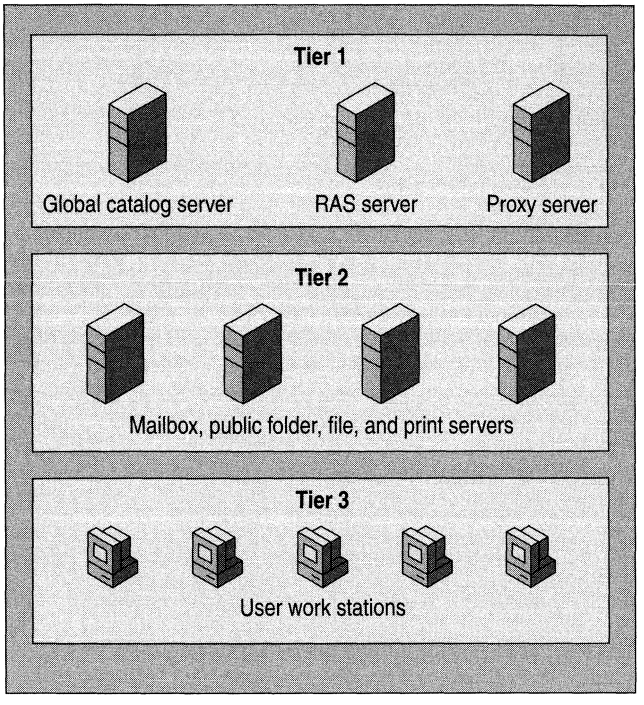


Figure 13.2 Three-tier virus protection

Tier 1: Backbone Infrastructure

A typical topology consists of three distinct levels. The first level, the infrastructure backbone, generally provides e-mail message switching, the directory, and routing and proxy services. These services are also linked to external communications such as Simple Mail Transfer Protocol (SMTP) and X.400 e-mail connectivity, and Web browsing, as well as partner and remote access connectivity. In most cases, these servers do not store end-user data, and clients do not interact directly with them.

Gateway and firewall protection products act at the switch or gateway level and search the media for virus signature files in the bit stream. Most products that operate at this tier are loaded on SMTP backbone servers and provide a means to scan content, block senders, and scan for key words. They provide a quick and unobtrusive mechanism by which Internet e-mail messages can be scanned and quickly quarantined or fixed. Although vendors of these products assert that they can reside on the same server as Exchange, it is a good practice to deploy gateway and firewall solutions on dedicated hardware.

Ideally, these products should receive automatic updates of virus signature files. To be effective the software needs to be kept as up-to-date as possible. In addition, products deployed at this level must not have a negative impact on gateway performance. Keep in mind that one server at this level might handle a load equivalent to 10 mailbox servers.

Virus protection at the backbone infrastructure level should encompass a number of services, including:

- E-mail (SMTP and X.400 gateways)
- HTTP, File Transfer Protocol (FTP), and any other external file transfer mechanisms
- Automated update of signature files

Products at this level include InterScan VirusWall by Trend Micro Incorporated.

Tier 2: Local Servers

The second level of a standard topology generally provides local mailbox and knowledge management services, as well as file and print services. This level hosts the servers that clients access directly to retrieve, store, and send messages, print documents, and store files.

Virus protection at the middle level should encompass a number of services, including:

- Mailbox and public folder scanning utilities for Exchange that use the Exchange virus scanning API.
- File-based scanning utilities.
- A process, whether automated or not, to update signature files.
- Custom utilities that target specific viruses or security alerts.

Products at this level must be able to scan the Exchange mailbox and public folder stores without affecting the processes by which transactions are logged. If you deploy products that do not use the Exchange 5.5 SP3 API, server performance could be degraded or data could be corrupted.

Tier 3: Desktop

The desktop is the entry point for the majority of data. This is where clients interact with client-side applications to read, create, send, and store local messages and files. Desktops differ from the local and backbone servers by performing numerous tasks such as word processing, spreadsheet manipulation, database operations, and Internet browsing. The number of desktops in an organization can vary greatly, from a few to hundreds of thousands.

Virus protection at the desktop level should encompass:

- Real-time and scheduled scanning capabilities.
- A process to update signature files on all client computers.
- End-user awareness and education about viruses. It is very important that you make end users aware of virus threats, e-mail policies, and escalation resources.

Protection at this level requires the installation of local scanning agents, such as Cheyenne Inoculan by Computer Associates or Norton Antivirus for Windows by Symantec Corporation. The main requirement of these programs is the ability to enforce installation on the desktop and to automatically distribute signature file updates.

Virus Scanning Support in Exchange

Quality virus protection support for Exchange servers was not available prior to the development of the Microsoft Exchange 5.5 SP3 API. Some third-party products advertised protection for the information store in eariier versions of Exchange, but most of these products affected performance or actually caused problems with the integrity and consistency of the store. The Microsoft Exchange 5.5 SP3 virus scanning API solved this problem by introducing antivirus support.

Virus Scanning API in Exchange 5.5 SP3

The Exchange Server version 5.5 Service Pack 3 (SP3) virus scanning API is an interface built specifically for Exchange; it takes into account the single-instance storage of messages and attachments. Microsoft provides this interface with Exchange 5.5 SP3 and with Exchange 2000. Only the virus-scanning interface is provided; you must obtain specific virus scanning solutions from third-party vendors specializing in virus scanning and removal.

The purpose of this API is to enable antivirus software developers to create simple, efficient means of detecting and cleaning harmful messages in the messaging infrastructure. All attachments can be scanned and intercepted as they enter the e-mail system and before they reach e-mail clients. The scanner examines only message attachments; individual text messages without attachments are not scanned.

Installation

The installation of the dynamic link library (DLL) that makes use of this new API is the responsibility of the vendors developing antivirus solutions. The Microsoft Web Storage System (also called the Information Store service in the Services snap-in) accesses the DLL by means of the Microsoft Windows® 2000 registry. When the Web Storage System starts, it scans the registry for the virus scanner information and, if it finds it, loads the DLL into memory.

After the DLL is loaded, the Web Storage System maintains a check timer on the scanner registry path to detect any changes that might be needed. The vendor installation program is responsible for proper versioning checks and for any parameter information that might be passed to the scanning DLL.

Because each server operates independently, you must install the virus-scanning DLL on every server that contains mailboxes. Servers that function only in a bridgehead or gateway capacity, and that do not contain mailbox stores or public folder stores, do not need to have the DLL installed.

Operation

Virus scanning functionality works because Exchange loads and keeps the DLL in memory as long as the Web Storage System is running. When a message with an attachment enters the Web Storage System, the attachment is stored in an attachment table; this is how Exchange maintains single-instance storage within the database. The virus-scanning DLL scans all the unscanned attachments in the table. When a virus is detected in an attachment, the message is marked as non-retrievable; calls from the client to open the message fail. If the attachment does not contain a virus, it is marked as clean and no further action is taken. If the client attempts to access the message while the attachment scan is in progress, calls from the client to open the message are delayed.

Messages marked as non-retrievable due to virus contamination are not necessarily lost forever. You can update the virus scanner with newer signatures and then rescan the messages to clean them. If a new virus is found and the message fails, the virus scanner cleans the infected messages and makes them accessible to users. Other approaches for thwarting the spread of viruses are also possible. For example, the virus scanner can change attachment extensions to make it difficult for programs to open files. The approach differs among antivirus software packages.

Because you must install the virus-scanning DLL on every mailbox server, redundant scans could occur on messages that are passed from user to user and server to server. However, when the DLL scans a message, it gives the message a new property to signify whether it is clean or infected. This property is then replicated to the other servers. When a message is forwarded to a different server, the scanner on that server sees that the message is clean and scanning does not occur. However, if the replicated server is running another vendor's scanning software or a more recent version of the existing product, the message is rescanned.

When the Web Storage System starts, it detects the virus-scanning DLL and places the DLL in a ready state. The DLL is a self-sufficient module that operates independently of the Web Storage System; it maintains a link to the attachment table and scans through all attachments in the background. The scanner also maintains the version keys in the registry and downloads the most recent version of the signature file. Since the virus-scanning DLL is a service of the Web Storage System, a call to close the Web Storage System also initiates a shutdown of the virus-scanning DLL. The vendor module must clean up any tasks after the shutdown call is initiated.

Administration

Microsoft does not provide an interface to administer the virus-scanning DLL. The vendor who develops the tool must provide the user with the interface needed to implement and use the service. These administrative options could easily be written as MMC snap-ins for System Manager.

Performance

You can expect some performance degradation after you implement the virus-scanning DLL. Because it scans every single message attachment that enters the server, the effect on performance depends on user load and message traffic flow rates.

The majority of processor time required by the scanner involves scanning attachments. The processing speed depends on the scanning software. Although the virus-scanning DLL runs in the Exchange Web Storage System. it is nevertheless an independent module; the rpeed and efficiency of the scanning process depends largely on the efficiency of the DLL itself. Microsoft will continue to work with vendors to make sure their code is as fast and reliable a5 possible. In addition, the primary task of the Exchange Web Storage System is to satisfy the needs of the mail system users. Great care has been taken by Microsoft to ensure that this new API does not, in any way, change the features and functionality of Exchange 2000 Server.

Logging Features

Exchange 2000 does not support the ability to log sender and recipient information for messages containing viruses.

Virus Scanning Support in Exchange 2000 Server

Exchange 2000 Server uses a similar virus scanning API to the one in Exchange Server 5.5. In addition, Exchange 2000 must be able to scan native MIME content. The main goal of Exchange 2000 is to make sure that the virus scanning API is backwards compatible. With changes in Exchange 2000 architecture, some new features may be added to the API to provide access to message properties and improve logging features.

Protection Without the Virus Scanning API

Most vendors do not yet have products that are compatible with the Exchange 2000 implementation of the virus scanning API. It may be necessary to ensure that adequate protection is supplied at the SMTP mail gateway and user desktop. Fortunately, many recent viruses have used a common file name for the infected attachment, which makes these viruses easier for users to recognize without the aid of antivirus software.

High Performance with the Virus Scanning API

The Exchange 5.5 SP3 virus scanning API enables antivirus software developers to obtain performance that is not possible with a MAPI notification approach, because the scanning DLL runs as part of the Web Storage System's store.exe process, in the client path. This is a different approach from asynchronous MAPI notification.

Because an antivirus software product must scan every new and modified item in the Web Storage System, running outside of this process consumes considerable processing power. Performance improves if the antivirus software product can perform simple checks on a large percentage of items inside the Web Storage System, rather than outside it.

Encrypted Messages

Encrypted messages cannot be scanned while they are encrypted

Virus Scanning Support in Outlook

Microsoft has recently added a feature to Microsoft Office 2000 that allows users to digitally sign macros. An organization can configure end-user desktops to execute only those macros that have been developed by the organization. This eliminates some of the risks posed by macro viruses. For more information about Office, see http://www.microsoft.com.

Microsoft has also posted an Office update that provides additional security for opening e-mail attachments. This update changes the attachment dialog box when certain attachments are opened. Users now see more explicit warning language and are required to save the attachment to the file system before opening it. This patch can help avoid accidentally releasing viruses that hide in executable files. Note that saving the attachment to the file system does not automatically remove any viruses that are present.

For the latest Outlook 2000, Outlook 98, and Outlook 97 Email Attachment Security Update, see http://officeupdate.rnicrosoft.com. You should continue to update your system with security updates as they become available. In addition, you should examine browser security settings, such as Microsoft ActiveX® and Java enabling. Having these features enabled can allow some types of viruses to infect systems more easily.

Spam Support in Outlook

In the same way that retailers and businesses use lists of postal addresses to send catalogs and other information to potential clients, businesses are increasingly using e-mail as a direct marketing tool. If you do not wish to receive these unsolicited messages, Outlook can search for phrases commonly used and automatically move such messages from your Inbox to a specified folder. The list of terms that Outlook uses to filter suspected spam can be found in a file called Filters.txt.

You can also filter messages by using a list of e-mail addresses flagged as undesirable. Outlook creates a list of such addresses; you can update it as unwanted messages arrive.

You can also install third-party filters that are regularly updated. For more information about filters that use updated lists of commercial and adult content senders, see the Outlook Web site at http://www.microsoft.com.

Virus Solution Vendors

You can use the information in this section to evaluate antivirus solution vendors and antivirus software products. Tables 13.1 and 13.2 provide checklists for evaluating antivirus software products.

Table 13.1 Antivirus product feature checklist

Things to Consider and Evaluate	Product 1	Product 2	Product 3
Does the vendor use the Exchange 5.5 SP3 API set?			
Does the virus-scanning DLL allow for quick and automated updates?			
Is the virus-scanning DLL written to ensure efficient performance?			

Table 13.2 Antivirus product measurement checklist

Things to Measure	Product 1	Product 2	Product 3
Run System Monitor and test and measure processor utilization by the DLL and Web Storage System.			
Test the auto-update feature of the DLL, if available.			
Test the DLL's ability to scan messages only once.			
Test the product's ability to prevent viral intrusion.			

After you narrow down the product selection, you can do a performance evaluation and a comparison of products.

Perform the following tests on each product:

- Attempt to pass a virus file through detection.
- Measure performance impact.
- Assess the usability of program features.

Vendors can provide disks containing viruses for testing purposes. These disks should be used only in a closed lab environment. You might also want to test the virus scanning software while the server is under stress or load. Several products, such as the Microsoft Exchange Server Load Simulation Tool (LoadSim.exe), are available for simulating a stressed or high-load environment.

The following table lists some common attachment types and associated viruses.

Table 13.3 Common attachments with viruses

Types of Viruses	Description
Executable virus	TROJ_EXPLOREZIP (also known as Worm.EXPLOREZIP)—an e-mail worm containing a damaging payload.
Known macro virus	The W97M_Empirical Word Macro virus.
Disguised known virus	Known virus signature—renamed as *.jpg to avoid detection if only scanning on .exe and not .jpg.
Archived and nested virus	Worm.Explorer virus—zipped and placed in an Outlook message file (.msg), which is then attached to the outgoing message.
Zero-byte .com file	Zero-byte file with .com extension—designed to test a known issue with some scanning programs, reported in August of 1999, where a zero-byte file could cause the scanner to stop responding.
Zero byte .zip file	Zero-byte file with .zip extension—designed to test a known issue with some scanning programs, reported in August of 1999, where a zero byte file could cause the scanner to stop working.

Detection Circumvention Tests

The following tests were devised to measure detection rates and attempts to defeat scanning detection.

Table 13.4 Virus detection tests

Test Message Condition	Expected Test Result
Message contains digital signature	Unable to detect without affecting signature. Ideally: Integrate with Certificate Server to scan mail.
Message is signed and encrypted	Unable to detect, encrypted. Ideally: Integrate with Certificate Server to scan mail.

Table 13.4 Virus detection tests *(continued)*

Test Message Condition	Expected Test Result
Message sent to uninitialized mailbox	Ideally: Initialize mailbox to scan incoming mail, otherwise first time logon to mailbox will allow access to virus.
Message sent with delay	Detection (immediately or on send).
Message sent with invalid return address	No problems (expect NDR if notifications to sender).
Message embedded in Outlook form	Detection.
Message sent to distribution list address	Detection except in the case of outbound SMTP (Custom Recipient) member of distribution list.
Message sent to public folder via post	Detection.
Message sent to public folder via an SMTP address	Detection.
Drag and drop file to public folder	Detection.
Message sent with private setting	Detection.
Attempt to catch virus while AV Service stopped or starting	Detection.
Message saved in Sent Items	Cleaned copy of message sent should be kept; but no active virus should be kept.
Message sent to .pst file as delivery location (Client logged on)	Detection.
Sent with invalid inbound address (creates NDR, but should not scan attachment)	Return without Scanning Attachment; NDR message should come back with virus intact. (Note: Consideration when under attack during viral breakout such as "Melissa.")
Sent with invalid Address (create NDR) with valid CC	Detection.

Performance Objects

Several virus scanning products add performance counters to the Performance console in Microsoft Windows 2000. These performance counters are used for scanning purposes and for detecting performance rates. You should also examine the standard performance counters in System Monitor. Counters to watch include processor time and memory used by the virus scanning process, message delivery time, disk usage and message queue length.

Monitoring Service

You should add the virus scanning service to existing monitoring systems, such as System Manager. You should also configure notifications so that you are alerted if the scanning service is offline.

Alert Monitoring

Link Crawling, a quick way to receive updates or alerts about potential virus threats, is part of Microsoft Internet Explorer 4 or later. The service is simple. After you subscribe to a Web site, Internet Explorer returns to the site on a scheduled basis and checks for updates. If an update is found, Internet Explorer can send an e-mail alert to the specified address.

To configure Link Crawling

1. Right-click the Active Desktop, and then click **Properties.**
2. Click the **Web** tab, and then click **New** to open the New Active Desktop Item wizard.
3. Type the URL you want to monitor (for example: http://www.cert.org/nav/whatsnew.htrnl), and then click OK.
4. Click OK to return to the **Web** tab.
5. Select the Web site you entered in step 3, and then click **Properties.**
6. Click the **Schedule** tab, and then configure a schedule to check for changes.
7. Click the **Download** tab, click **When this page changes, send e-mail to,** and then configure an e-mail account.

For the most up-to-date links about viruses and virus technology, search the Internet using terms and concepts from this chapter.

CHAPTER 14

Server Availability

Martyn Davis, Consultant, Microsoft

Providing a high-availability solution for a system involves many areas of an organization's technical and operational infrastructure. This discussion focuses on planning a high-availability infrastructure for Microsoft Exchange 2000 Server.

Many of the recommendations in this discussion apply to large Exchange 2000 deployments and may not apply to many smaller installations. However, a high-availability and resilient architecture may also be valuable in smaller installations if you use Exchange 2000 Server as a platform for workflow, collaboration, and line-of-business applications.

Detailed configuration of the hardware used for high-availability architectures is not included in this chapter. For information about hardware and hardware configuration, see the software and hardware product documentation. Some vendors include specific information about using Exchange on their own hardware products.

In This Chapter

Defining High Availability Architecture

Determining Availability Requirements

Implementing Availability Technologies

Creating Availability Processes

Defining High Availability Architecture

When a system operating in your company fails, the consequences can be devastating. Lost revenue, interruption in service to customers, and an idle workforce are all possible outcomes. Whatever the situation, recovering from a failure incurs costs in many areas of your organization. Other hidden costs may include the loss of good will; damage to your reputation among customers. suppliers, and partners: and the perception that your organization is not equipped to satisfy customer needs. Therefore, when you determine your requirements for high-availability, consider your organization's total cost for system failures.

The following terms are useful for discussions about availability:

Availability The proportion of time a system is functioning. Availability is often described by a percentage. Table 14.1 shows common availability figures.

Table 14.1 Availability figures

Availability	Downtime per year
99 percent	3 days, 15 hours, 21 minutes
99.9 percent	8 hours, 44 minutes
99.99 percent	52 minutes
99.999 percent	5 minutes

Mission-critical Those information systems that are essential to the function and success of an organization. Many organizations treat messaging as a mission-critical application. Exchange 2000 Server is particularly important if it provides a platform for collaboration, workflow, knowledge management, and line-of-business applications.

Mean time between failure A statistically derived length of time a user may reasonably expect a component, device, or system to work between two incapacitating failures.

Reliability A measure of how dependable a system is. Reliability can also be considered the combination of availability and data integrity.

Recovery point The amount of data that is lost because of a system failure. A possible recovery point objective for organizations is to provide a recovery point of zero (no data loss).

Recovery time The amount of time needed to return a system to operation. The recovery time objective is to reduce the recovery time for a particular type of failure.

Determining Availability Requirements

Before investing time and money to design and deploy a high-availability solution for Exchange, take the following actions:

- Determine the cost of downtime.
- Define the recovery point objective and recovery time objective.
- Identify and understand the events that can have a negative impact on both Exchange 2000 Server and the services and systems that rely upon it.
- Understand the vulnerability of the components, processes, and systems that rely on Exchange 2000 Server.

After you consider these factors, determine which technology and processes you can use to achieve the availability you need.

After you complete these actions, you will have an understanding of the following:

- The financial costs and impact of a failure within your Exchange organization
- The expected recovery time and recovery point if a disaster occurs
- The risks and events particular to your location and business that can cause a failure of Exchange
- Other software components within your organization and within the Exchange infrastructure that can cause failure, and the types of failure that they can cause
- The technology and processes that you can use to provide high-availability for Exchange

The following sections describe how to determine your own availability requirements and implement the necessary technology and processes.

Cost of Downtime

What happens to your organization if your internal telephone system fails? What impact would this have on the day-to-day business operations? Communications between individuals, departments, customers, and partners can be affected.

Many organizations that use e-mail underestimate how vital messaging service is to their internal and external communications network. Because e-mail is also frequently used for routing in workflow applications, the services provided and business functions supported by e-mail services can be far-reaching.

If you use Exchange 2000 for collaborative and knowledge-management applications, secure and reliable storage for Web sites, and document and multimedia data repositories, many areas of your organization rely on Exchange 2000 for more than e-mail.

To quantify the level of availability required for Exchange 2000, ask who or what is adversely affected when:

- **A mailbox server fails**

  Communication between individuals, groups, and external contacts cease. Information stored on the server is not available. Operational processes may stall because managers cannot carry out authorization, verification, and approval of business processes. Personal productivity, workflow, collaboration, knowledge management, and line-of-business applications that rely on user interaction or message storage are affected.

- **A public folder sewer fails**

  Critical business information coming into your organization, for example, news feeds, cease. Applications, document repositories, Web sites, and knowledge management systems stop working. Personal productivity suffers.

- **A connector server fails**

  Message delivery between internal and external systems stops. Virus and content scanning may not occur. Business processes are interrupted. External customers cannot communicate with your organization; sales and revenue suffer.

- **A network failure occurs**

  Depending on what type of failure occurs, and where it occurs, different things can happen. Message delivery between Exchange 2000 components slows down or stops. Unreliable access from clients and servers to Active Directory directory service may cause logon failures. Communications to customers and partners may be affected.

- **One or more Active Directory global catalog servers fail**

  Global catalog servers act as the global address list for both clients and Exchange 2000 servers. More than 90 percent of all interaction with Active Directory occurs with the global catalog. Without access to these servers, the corporate address list is not available. Users may not be able to log on to Microsoft Windows 2000.

- **One or more Active Directory domain controllers fail**

  Exchange 2000 uses Active Directory to store configuration information. Without Active Directory, it is not possible to manage Exchange; queries for directory information by Exchange 2000 servers can fail, producing communication errors and a decrease in overall stability.

It is evident that Exchange 2000 can support many applications, processes, and interfaces and is key to the personal productivity of many of your employees.

Losses can be measured many ways, but if money is the measure, then the figures can be astounding. In a report called "International DARTS '98," The Standish Group states that costs of downtime typically range from $1,000 to $27,000 per minute. Furthermore, they report that in some cases, the cost of downtime for a single incident has exceeded $10,000,000.

System Vulnerability

Because of the large marketplace of independent vendors that provide hardware components for integration into workstations and servers, there is always the chance for incompatibility between components. Similarly, although vendors are working at great lengths to improve mean time between failures, it is inevitable that failures occur.

It is important to understand the vulnerability of your environment to failure, with respect to key focus points.

- **Hardware** This relates to the servers, and the components within them, that are to support Exchange 2000 services. Typical considerations are uninterruptible power supply, disk subsystem performance and reliability, and processor failure.
- **Network** Networking components from cabling to complex routers and multiplexers are essential in providing a reliable service for Exchange users. A robust network infrastructure is necessary to build a high-availability solution for Exchange 2000.
- **Software** Because modern application software is so complex, it invariably causes some problems. These problems can arise from an individual application or the interaction between two or more applications. A key indicator of your vulnerability in this area is how well are you able to resolve software errors or conflicts with the application of hot fixes, updates, and service packs.
- **Management Process** How well do your service management processes and procedures support recovery work in the instance of failure? Are elements such as configuration management, capacity planning, and change control defined and understood sufficiently to provide the safety net for disaster recovery and the ongoing operation of Exchange 2000?

Recovery Point and Recovery Time

The degree of availability for Exchange 2000 depends on your organization's needs. This section examines availability in terms of recovery point and recovery time.

The recovery point focuses on the data that must be restored in the event of a failure. The recovery time determines the amount of time required to return the affected Exchange 2000 system to operation. There is trade-off between time and data, which you must consider when determining the availability of your Exchange 2000 environment.

You can distribute many of the components of an Exchange 2000 infrastructure, including connector servers, global catalog servers, and mailbox servers, over a wide geographical area. Each of these components can provide dedicated or multiple functions as part of the whole environment. It is likely that different levels of availability will be required for different aspects of Exchange 2000.

For example, message routing can be handled by a series of Exchange 2000 servers, thereby providing alternative routes for message delivery. If alternate routes have been built into the architecture, when one of these servers fails, message flow can continue. Therefore, the recovery time for this system can be longer than for a mailbox server. Similarly, the recovery point would not be critical because the data on that server is primarily queued messages.

There are several means to reduce recovery time and recovery point within Exchange 2000:

- Multiple databases and storage groups
- Resilient message routing
- Data replication
- Simultaneous online backup and simultaneous online restoration

Causes of Downtime

After you look at the cost of a failure for your organization and identify the availability requirements for your Exchange 2000 environment, it is important to understand the events that can prevent you from keeping the system in operation.

- **Planned downtime** This downtime allows you to perform hardware upgrades, software upgrades, and configuration changes. If you use clustering to perform "rolling upgrades," you can minimize downtime.
- **Component failure due to hardware or software issues** Vendors have been working to improve the mean time between failure rates for individual components. However. the complexity of today's hardware and software components still produces failures.
- **Human intervention** There are many circumstances when human actions, either malicious or unintentional, can have a negative impact on your environment. Planned administration can even lead to problems. For example, physically moving a server between locations may cause a hardware failure when that server is restarted.
- **Building-level incidents** These include disasters that can affect a room, building, or an entire location. Problems such as fire, flood, power outages, or political activities can interrupt service by physically damaging or destroying hardware, terminating power to essential computers, and preventing access to systems.
- **Metropolitan area disaster** Disasters such as fire, flood, blackouts, and riots can affect entire towns and cities, potentially affecting systems located in the vicinity.
- **Regional events** Even wider events can affect the operation of your Exchange 2000 environment. These events include earthquakes, hurricanes, and military and political upheavals that can impact your operations over hundreds of square miles.

A recent report by Gartner, "Making Smart Investments to Reduce Unplanned Downtime," found that about 80 percent of all unplanned outages are caused by application errors or operator errors; the remaining failures result from hardware, operating systems, or environmental issues. Operator errors include incorrectly applying procedures in a given circumstance, performing a task incorrectly, or not performing a task. Application failures include errors in software, performance issues, and incompatibility problems.

Implementing Availability Technologies

This section highlights some of the technologies that can be applied when developing a high-availability architecture. There are several architectural and component-level design considerations that you can use to improve resilience and introduce redundancy into the environment

Server Hardware

- **Power supply** There are several approaches to ensuring that an adequate power supply is provided. Many server configurations support redundant power supply units within the server, allowing the secondary unit to take control in the case of failure. Providing an uninterruptible power supply is another common method of protecting electrical input. A solution for small single server can provide backup power ranging from minutes to a few hours, thereby allowing sufficient time for a server to shut down in a controlled manner. An uninterruptible power supply can help prevent potential data loss or database corruption. For maximum protection, as is typically desired for larger data centers or computer rooms, you can use backup generators.
- **Network interface** Having a robust and resilient server is worthless if no users or servers can access it. You can improve a high-availability architecture by using two network interface cards: reserve one as a ready spare or for connecting to two different LAN segments. You can also design and deploy a high-availability network infrastructure with redundant links and intelligent routing.

Distributed Services

- **Distributed environment** Typically Exchange **2000** will be installed on a number of servers throughout your organization's infrastructure. In environments with multiple physical locations, making use of routing groups and placing connector servers to reduce single points of failure improves the overall availability of the system. Similarly, public folder stores can be distributed over a wide geographical area, not only to improve performance for user and application access, but as another method of reducing data loss in a localized emergency or hardware failure.
- **Data replication** One method of backing up data is that you can keep a copy of information in other locations. Public folder replication allows public folder stores to replicate to other servers in the organization. This improves response time for users and applications and provides a backup of public folder data.

Multiple Databases

Exchange 2000 supports the implementation of multiple databases and introduces some new design concepts for the specification of Exchange 2000 servers. This is a major enhancement from Microsoft Exchange 5.5, which supports only one private and one public database for each server. The following data storage structures are included in Exchange 2000.

- **Storage groups** A storage group is a collection of Exchange databases that share a set of transaction logs. A maximum of five databases per storage group can be made up of any combination of mailbox and public folder stores. A single node running Exchange can support up to four storage groups; in a cluster, the total number of storage groups on a single node must be four or less, regardless of how many Exchange virtual servers exist on the node.
- **Database** You can define an Exchange database for use as either a public folder store or a mailbox store. A user's mailbox is assigned to a mailbox store. You can back up, restore, mount, and dismount databases independently of each other. Exchange supports up to five databases in a storage group.

Exchange 5.5 can support one private information store and one public information store on a server. This means that all mailboxes on a single server must be in a single database. With the removal of the 16-gigabyte (GB) limit in Exchange 5.5, Exchange 5.5 servers can host several thousand users, requiring 30 GB, 40 GB, 50 GB or even more disk storage. This results in:

- Potentially long backup times
- Long restoration time
- High impact to all users in the event of corruption or disk failure

In many cases, after a problem has surfaced, you must take the database offline and defragment the database to fix it. The defragmentation can take more than a day for a large database. You do not need to restore the database.

The most common method for providing resilience and increasing fault tolerance for data storage in a single machine is to implement a redundant array of independent disks (RAID) solution. This includes simple mirroring between two disks (RAID-1) and mirrored stripes or striped mirrors (RAID-0+1 and RAID-5). For more information about RAID, see "Server Sizing" in this book.

Even with these improvements, there is no reduction in the storage requirements on the Exchange 2000 server. Nonetheless, storage groups and multiple databases allow for more careful division of data. Table 14.2 shows how storage groups and databases can be configured.

Table 14.2 Storage group configurations

Storage Group	Mailbox Store	Users	Maximum Store Size
SG1	DB1	220	7.7 GB
SG1	DB2	250	8.75 GB
SG1	DB3	250	8.75 GB
SG1	DB4	250	8.75 GB
SG1	DB5	250	8.75 GB
SG2	DB1	250	8.75 GB
SG2	DB2	250	8.75 GB
SG2	DB3	250	8.75 GB
SG2	DB4	250	8.75 GB
SG2	DB5	250	8.75 GB
EXECSG	EXECDB1	30	1 GB

Case Study: Balancing Databases

Lakes & Sons have a single Exchange 5.5 server servicing 2,500 user mailboxes; another server in the organization handles public folders. Each mailbox has a limit of 35 MB of data, thereby making a maximum mailbox store size of 87.5 GB of data (this assumption does not account for disk space savings based on single-instance message storage). Therefore, every backup and restore handles 87.5 GB of data; all 2,500 users will be affected if a corruption occurs in the database and offline recovery utilities are used, or if the database must be recovered in its entirety.

When Lakes & Sons implement Exchange 2000, they split the 2,500-user load across a number of databases. They decided that a 20 GB maximum for a single database was the goal; however, a more important factor in the decision to split the load was to reduce the impact to users in the case of database corruption. They decided on a 300-user maximum for each database. Table 14.3 describes how Lakes & Sons achieved the user split.

The split reduces impact to a maximum of 300 users in the event of a single database corruption (all other users are still able to use the system). The split also reduces backup time by implementing a parallel backup strategy. In addition, the split allows you to increase the mailbox size limit to 60 MB per mailbox and the mailbox store size limit to 20 GB.

Intelligent Routing

Providing a high-availability Exchange 2000 infrastructure also requires a focus on message routing. Because an Exchange 2000 organization adds no value if messages cannot be delivered, defining and deploying a robust message routing topology is critical.

The placement of routing groups, routing group connectors, and connector servers is important for eliminating as many single points of failure as possible. The following items describe the key elements of message routing:

Routing group A logical group of Exchange 2000 servers with good network connectivity between them. These servers communicate directly with one another by using Simple Mail Transfer Protocol (SMTP) for message delivery.

Routing Group connector An SMTP connection configured for one or more Exchange 2000 servers in one routing group, and one or more Exchange 2000 servers in another routing group.

Exchange 2000 uses routing groups and routing group connectors to implement the Open Shortest Path First (OSPF) protocol to calculate message routing paths. Therefore, it is possible to develop complex and robust routing topologies for your organization.

OSPF is commonly used at the networks layer. This means that when a network link fails, the network seeks an alternate route for network traffic. When OSPF is implemented in the network layer, Exchange 2000 is not aware of the routing adjustment. However, you cannot fix the failure of a network card or connector server at the network layer; in these cases, correctly configured Exchange 2000 routing groups mitigate the impact.

Routing groups provide redundancy when used in a many-to-many relationship as shown in Figure 14.1.

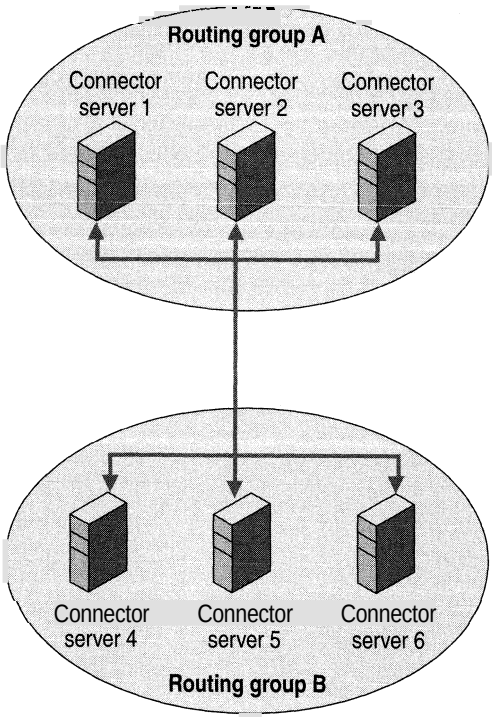


Figure 14.1 Routing group many-to-many relationship

In the illustration, a message on connector server 1 traveling from routing group A to routing group B attempts to connect to connector server 4 in routing group B. If for any reason connector server 4 is unavailable, connector server 1 will try to communicate with another connector server, for example, connector server 5, to deliver the message.

Another important consideration when providing a resilient routing topology is the placement and role of connector servers. Locally configuring multiple routing groups and connector servers without considering the underlying physical network can negate the benefits of routing groups. Similarly, using only one connector server for multiple routing groups can also reduce the number of benefits that this technology can offer. Consider the relationship between routing topology and the physical network in Figure 14.2.

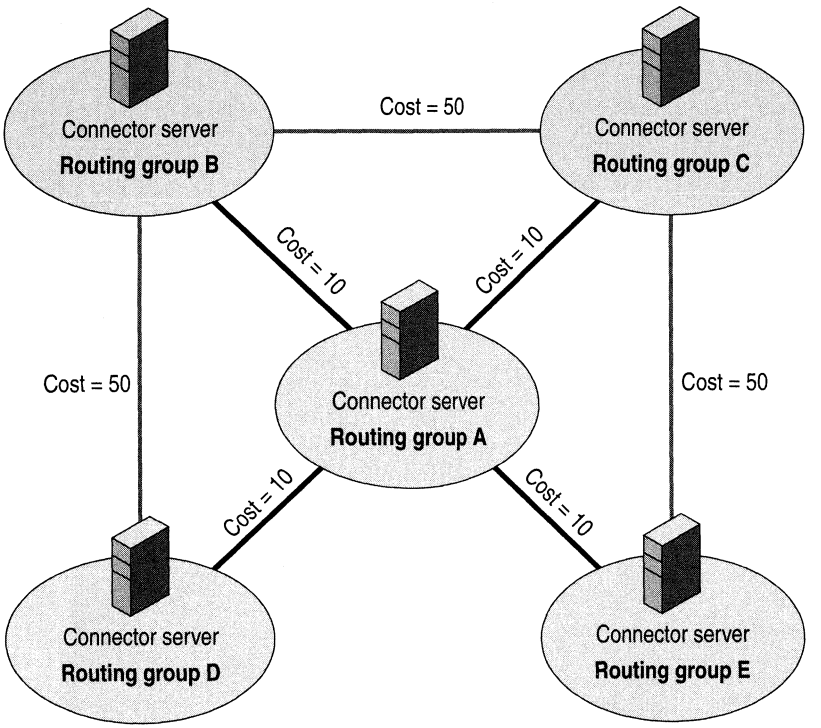


Figure 14.2 Routing group resilience

In the illustration, a message sent from routing group D to routing group E travels through the central routing group A. If the connector server in routing group A fails, messages travel between routing groups D and E through routing groups B and C. However, if the physical network between routing groups A and D is the same as between routing groups B and D, the alternate route is also unavailable to Exchange 2000.

Server Clusters

Exchange 2000 can balance access to all servers (also called nodes) in a server cluster, thereby making more efficient use of hardware. When a failure occurs in a server cluster that runs Exchange 2000, the Exchange virtual server fails over to another node. Each node starts with one virtual server that has one or more storage groups. When a node fails, the responsibilities for the virtual server on that node are transferred to another node in the server cluster.

A server cluster for Exchange consists of one or two server nodes, each with two network cards, an internal disk subsystem for the operating system and application executables, and a small computer system interface (SCSI) interface attached to one or more shared storage devices.

When Exchange 2000 is installed on a Windows 2000 cluster, virtual servers are configured for each node. Each virtual server is responsible for one or more storage groups. Keep in mind that a node or stand-alone server can operate with a maximum of four storage groups; therefore you must ensure that in the event of a failover, any remaining nodes are not supporting more than four storage groups.

> **Note** At this time. Windows 2000 Advanced Server and Exchange 2000 support no more tnan two nodes in a server cluster.

Table 14.3 provides examples of virtual server and storage group configurations for a two-node server cluster.

Table 14.3 Storage group configurations

Exchange Virtual Server Names	Number of Storage Groups
EVS1 EVS2	1 1
EVS1 EVS2	2 2
EVS1 EVS2	1 2
EVS1 EVS2	1 3

Using an Exchange 2000 cluster provides your organization with the ability to deploy a high-availability solution for mailbox and public folder data by using industry-standard hardware components.

Front-End and Back-End Architecture

In addition to providing a resilient architecture for routing, public folder stores, and mailbox stores, Exchange 2000 provides an approach for separating client access from client data for non-MAPI clients, such as Outlook Web Access, Internet Message Access Protocol (IMAP), and Post Office Protocol version 3 (POP3) clients. You use this service by deploying a front-end and back-end server architecture. This essentially defines a set of Exchange servers for data and a set of Exchange servers as protocol servers that service HTTP, IMAP, and POP3 protocols. MAPI clients must connect directly to a back-end mailbox server.

The main benefits of a front-end and back-end architecture are a unified namespace, the ability to isolate mailbox servers, and reduced overhead for Secure Sockets Layer (SSL) encryption.

- **Unified namespace** A unified namespace provides easier administration of multiple servers. For example, if you have three computers running Exchange, you typically divide the user load by configuring certain users to connect to Londonl, others to London2, and the rest to London3. If all servers are part of a front-end and back-end architecture, a single server name provides user access to multiple servers in your organization. You can configure clients to connect using the name LondonServers; software or hardware load balancing randomly distributes the load among the three front-end servers. The front-end server queries Active Directory, and sends the request to the correct back-end server. Also, when you want to move mailboxes from one server to another, clients do not need to reconfigure the name of the server to which they log. As your user population grows, you can add front-end servers without reconfiguring clients.
- **Isolate back-end servers from malicious attacks** A front-end and back-end architecture allows you to deploy the system in a secure manner. Back-end servers can reside in a private, isolated network (behind a firewall, or on a different subnet). Front-end servers can have two network cards and reside in both the external network and the secure, internal LAN. You can configure the front-end server by using TCP/IP Filtering to listen only on the designated protocol ports, thus limiting the opportunity for malicious attacks.
- **Reduced overhead for SSL** SSL connections require encryption and decryption steps that are processor-intensive and can affect performance. If your servers are deployed in a front-end and back-end architecture, the front-end server can process the encryption with the client. When the front-end server and back-end server communicate, they can do so without the overhead of SSL encryption. When you consider the benefits of front-end servers combined with back-end server isolation, you create a secure and efficient topology.

Horizontal scalability and fault tolerance are added benefits of the front-end and back-end architecture. If the front-end servers are not able to keep up with the demand, it is easy to add a new front-end server and put it into production. Because the front-end servers act as protocol converters and do not store any data, you need to only install Exchange 2000 and update Domain Name System (DNS) to include the new server name. You can provide fault tolerance by using a common IP hostname for all of the front-end servers, and then let the round-robin mechanism of the DNS server provide a specific IP address for each request. Because the front-end servers do not save any state information, another front-end server can handle the request if the first one fails.

When a client, such as Outlook Web Access, contacts a front-end server, that server queries Active Directory to determine the user's home server and database. The request is then passed on to the appropriate back-end server and a session is established. In this case, the client communicates through HTTP to the front-end server, which in turn uses HTTP to communicate to the back-end server.

By using Network Load Balancing with a set of front-end servers, it is possible for your clients to always connect in the same way, regardless of the Exchange server their inbox resides on. This adds fault tolerance when accessing user data and provides separation between users and their mailboxes. You should note that Network Load Balancing does not verify that the software has failed, only that the hardware has failed. This can allow some connections to appear operational even when there is no access to data. In such cases, you can move a mailbox between different back-end servers without impacting the user.

Note Front-end and back-end architecture is best suited to non-MAPI clients such as Outlook Web Access. If you use a MAPI client in this configuration, direct connection between the Outlook (MAPI) client and the back-end server would take place.

Figure 14.3 combines many of the elements already discussed in this chapter. The combination of multiple storage groups and databases in a server cluster, deployment of a front-end and back-end architecture supported by Network Load Balancing, and public folder stores replicated across both server clusters provide a high-availability solution for Exchange 2000.

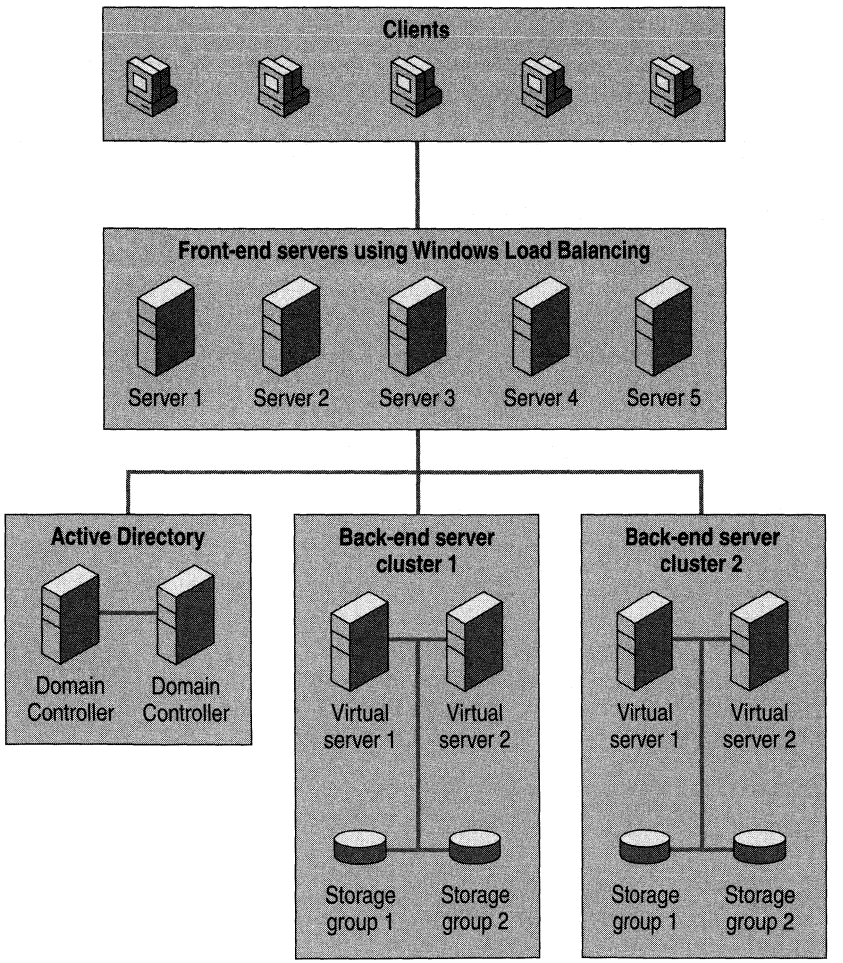


Figure 14.3 Scalable and resilient Exchange 2000 architecture

In addition to these technologies, the redundant and distributed nature of Active Directory reduces single points of failure in the system.

Creating Availability Processes

Even if you implement hardware and software solutions to improve availability, it is important to realize that these methods address the minority of causes of system downtime. Even if you invest in complex hardware architectures, you must allow time and provide resources for the design and testing of supporting processes and procedures.

Focus on the following key areas when managing high-availability architecture:

- **Training** Technical training for the operation of Windows 2000 and Exchange 2000, specific hardware training on the computer systems deployed to deliver Exchange, and training in operational and support procedures is necessary.
- **Change control** Understanding the impact that a hardware or software change has to your environment is essential. Developing processes, such as scheduling, impact analysis, approval, testing and deployment planning, is essential to introduce changes into a complex environment.
- **Backup and restore procedures** Determine the correct approach and schedule for backing up data, storing backup media, restoring from backup, and testing. Without these, recovery from disaster will be severely hampered.
- **Disaster recovery** Disaster recovery planning requires that you determine how to recover not just a single server, but potentially a large proportion of the Exchange 2000 environment as quickly as possible. Building-level and metropolitan-area events could cause the initiation of disaster recovery procedures. After you have carried out tasks such as network provisioning, hardware acquisition, and identification of critical business functions, you should regularly test disaster recovery plans.
- **Problem management** Defining processes to speed problem determination and resolution, including building an internal knowledge base that details previous support issues, will assist in the recovery of systems. Clearly defined interfaces with support organizations such as Microsoft Product Support Services and hardware vendor support are also critical.
- **Capacity planning** Being proactive in the measuring and reporting of the operational state of systems such as Exchange allows for planned and managed maintenance of hardware and software.

In summary, the technologies available today and those being introduced over the coming months and years will only be as effective as the processes and procedures that your organization develops, implements, and maintains.

Conclusion

Addressing the availability of a distributed application service such as Exchange 2000 can be challenging. Keep in mind that high availability for mailbox data and public folder data, and message routing are priorities in your Exchange 2000 organization.

The definition of a suitable high-availability architecture for Exchange 2000 in your organization will be achieved by comparing the costs, both financial and operational, of experiencing systems failure, against the investment costs and operational support costs associated with ownership of the environment. From these criteria, coupled with scalability, security, and performance requirernents. you can define a suitable architecture. However, it is essential that you determine the high-availability requirements and plan for all elements of Exchange 2000 prior to deployment.

CHAPTER 15

Server Sizing

Jens Trier Rasmussen, Senior Consultant, Microsoft

Microsoft Exchange 2000 Server is highly scalable—it expands easily to meet your needs as your company grows. There are two dimensions to scalability: horizontal scalability, in which you add more servers doing the same tasks, and vertical scalability, in which you configure a single server for better performance. This chapter focuses on vertical scalability.

These recommendations and best practices target an environment with a large number of users and servers. If your Exchange environment serves only a limited number of users on a small number of servers, these recommendations might not be relevant. However, if you anticipate that your company will grow in the future, you should consider implementing the recommendations and best practices that apply to your environment. It is always easier to configure the servers before they are in production; it is more difficult, complex, and risky to make changes in a production environment.

Microsoft Windows Clustering and performance optimization for Exchange 2000 components are also important factors in server sizing. For more information about clustering and performance tuning, see the Microsoft Windows 2000 Server documentation and Microsoft TechNet. Although several of the recommendations in this discussion also apply to Windows 2000 components such as the Active Directory directory service, WINS service, and DHCP service, this chapter does not discuss these components.

In This Chapter

Exchange 2000 Server Types

Processor Configuration

Disk Configuration

Memory Configuration

Network Configuration

Exchange 2000 Server Types

Exchange 2000 Server includes several different components—for example, the Microsoft Web Storage System, the message transfer agent (MTA), the connectors, which can all be run on the same server. However, to optimize each component for maximum performance it is recommended that you separate the different components, wherever possible, on dedicated servers. This has the added benefit of making the software configuration and interoperability of the servers less complex; thus, the servers are more stable and reliable.

To illustrate this point, consider a situation where you run Exchange, Microsoft SQL Server, and a third-party SQL application on the same server. The application requires a specific service pack that has not yet been tested with Exchange. Installing the service pack may cause stability issues with Exchange. If you have servers that are dedicated to Exchange, you avoid this issue.

You can have the following types of Exchange 2000 servers:

Mailbox servers This server type is dedicated to holding user mailboxes and servicing mailbox data to the clients by using MAPI, HTTP, Internet Message Access Protocol 4 (IMAP4), or Post Office Protocol 3 (POP3).

Public folder servers This server type is dedicated to holding public folders and servicing public folder data to clients by using MAPI, HTTP, HTTP-DAV, or IMAP4 protocols.

Connector servers This server type is dedicated to running the various connector types, for example, X.400, Simple Mail Transfer Protocol (SMTP), Lotus Notes, and GroupWise. If a specific connector is heavily used, you should dedicate a specific server to running this connector and then configure another server to run the other connectors.

Front-end servers This server type interacts directly with clients and passes requests from clients to back-end servers, which maintain the mailbox stores and public folder stores.

Data conferencing servers This server type is dedicated to running data conferences on Exchange 2000 Conferencing Server.

Video conferencing servers This server type is dedicated to running video conferences on Exchange 2000 Conferencing Server.

instant Messaging sewers This server type is dedicated to running the Instant Messaging service. You can have Instant Messaging home servers hosting users and Instant Messaging routers routing requests.

Chat servers This server type is dedicated to running the Chat service.

The data conferencing, video conferencing, Instant Messaging, and chat servers are often called real-time collaboration servers.

Whereas the directory replication bridgehead server and Outlook Web Access server are commonly used in Exchange 5.5 environments, they are not used with Exchange 2000. Active

Directory handles all directory services for Exchange 2000: because HTTP access is a built-in feature of the Web Storage System, there is no need for a dedicated Outlook Web Access server.

> **Note** Exchange 2000 servers installed in an Exchange 5.5 site use the Site Replication Service (SRS). The SRS is similar to the Exchange 5.5 directory service, although MAPI is disabled.

Exchange 2000 does not contain its own directory but uses Active Directory. To speed access to data, Active Directory uses the DS Access component. This component caches the result of queries to Active Directory in shared memory, which imposes storage requirements on the swap file. It is recommended that you add 64 MB to the paging file for the DS Access cache.

For best Active Directory performance, dedicated domain controller servers and global catalog servers are recommended. When designing your Active Directory infrastructure for Exchange 2000, remember that Microsoft Outlook users and Exchange 2000 server components make extensive use of the global catalog. It is generally recommended that you do not configure Exchange 2000 servers as Active Directory domain controllers or global catalog servers. This configuration requires large processor and memory resources. However, in certain situations, it might be cost effective to configure the Exchange server as a domain controller or global catalog server. An example of such a situation is a branch office where a local Exchange server serves a small number of users.

The remainder of this chapter provides recommendations for configuring the following components of various types of Exchange 2000 servers:

- Processor
- Disk
- Memory
- Network

Processor Configuration

The recommended CPU for Exchange 2000 is a 500-MHz Intel Pentium II Xenon processor. In general, the CPU should be the best possible at time of purchase; at this time this is a Pentium III processor with a speed greater than 500 MHz. In addition, it should have a 2-MB level 2 (L2) cache.

Capacity for Mailbox and Public Folder Servers

Mailbox and public folder servers can serve a large number of users who are active at any given moment. The requirement for processing power for the server depends on the number of users and the usage profile for those users. Table 15.1 lists the recommended number of CPUs for a user population.

Table 15.1 Recommended CPU configuration for mailbox servers

Number of Mailboxes	Number of Processors
Less than 500	1
Between 500 and 1,000	2
More than 1,000	4

Note that when determining the number of mailboxes you can have on a server, the important factor is usually not the raw processor capacity. The most important factor is normally the storage limits and the defined service levels for the maximum time required to perform a backup, recover from a disaster, restore a single database, restore a single mailbox, and restore a single document. Based on this maximum time, you can calculate the maximum database size and the number of mailboxes.

Service levels and response time are the primary factors to consider for public folder servers. Your requirements depend directly on how your Exchange organization uses public folders.

Capacity for Connector Servers

In large Exchange 2000 implementations you often see a routing infrastructure built around a hub-and-spoke configuration with one or more hubs. The spokes form routing groups and have dedicated routing group connectors running on dedicated connector servers. The routing group connector servers on the spoke communicate with the routing group connector servers in the hub. Depending on traffic volumes, the hub routing group connector server can also handle communication to the Internet and connection3 to other messaging systems.

In general, connector servers require high CPU performance to handle the processing and conversion of messages. The recommended CPU configuration of these different types of connector servers is shown in Table 15.2.

Table 15.2 Recommended CPU configuration for connector servers

Connector Server	Number of CPUs
Spoke routing group connector server	2
Hub routing group connector server	4

Capacity for Data and Video Conferencing Servers

In general, data and video conferencing servers require high CPU performance to handle the processing and maintenance of conferences. The recommended CPU configuration is shown in Table 15.3.

Table 15.3 Recommended CPU configuration for data and video conferencing servers

Simultaneous Users	Number of CPUs
Fewer than 500	1
Between 500 and 1,000	2

Capacity for Instant Messaging Servers

In general, Instant Messaging home servers require high CPU performance to handle the processing and maintenance of presence information. The recommended CPU configuration is shown in Table 15.4.

Table 15.4 Recommended CPU configuration for Instant Messaging sewers

Instant Messaging Users	Number of CPUs
Fewer than 5,000	1
Between 5,000 and 10,000	2

Capacity for Chat Servers

In general, Chat servers require high CPU performance in order to handle the processing and maintenance of statuses. The recommended CPU configuration is shown in Table 15.5.

Table 15.5 Recommended CPU configuration for Chat servers

Chat Users	Number of CPUs
Fewer than 10,000	2
Between 10,000 and 20,000	4

Disk Configuration

Disk input and output (I/O) performance is an important part of the total performance of an Exchange 2000 server. You should investigate the desired disk characteristics of each type of server, including the following elements:

- Placement of data
- Redundant array of independent disks (RAID) technology
- Disk controller cache settings
- Exchange 2000 paging file settings
- Capacity considerations

It is important that Exchange 2000 use different storage areas for different purposes. The storage should always be fault tolerant; that is, it should use RAID-5, RAID-0+1 or RAID-1 technology. To improve fault-tolerance, you should use redundant RAID controllers in active-active mode and have one or more standby disks per disk enclosure.

Figure 15.1 shows the strengths of different types of RAID technology. The closer a version of RAID is to a corner of the diagram, the more it exhibits the trait in that corner.

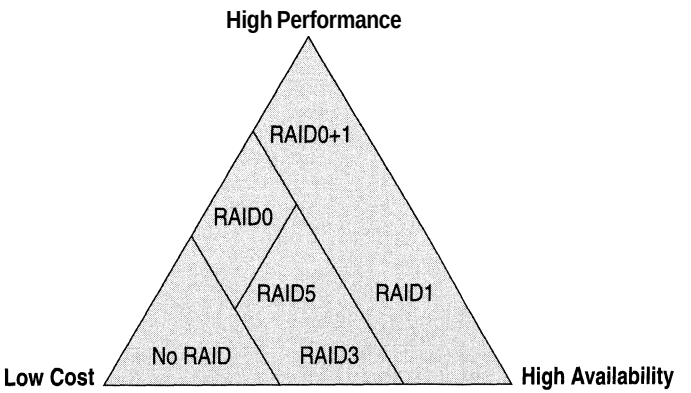


Figure 15.1 RAID performance, cost, and availability

For example, in Figure 15.1, RAID-1 and RAID-0+1 are in a segment of the diagram closest to the "high performance" and "high availability" corners; therefore, they exhibit high performance and availability, but not lower cost.

For more information about RAID, see *The Raidbook:* A *Handbook of Storage Systems Technology* by Paul Massiglia, 1997, Peer to Peer Communications, and the RAID Advisory Board at http://www.raid-advisory.com.

In the following discussion the term volume describes two or more physical disks grouped together using RAID-5, RAID-0+1, or RAID-1 technology. Note that this is different from two or more partitions of the same physical disk.

Disk Configuration for Mailbox and Public Folder Servers

With respect to disk I/O characteristics, mailbox and public folder servers behave similarly. Each mailbox store or public folder store consists of a rich-text .edb file (for normal messages) and an .stm file (for native Internet content). Mailbox and public folder stores exist within a storage group. Each storage group shares one set of transaction log files. You can have one or more storage groups on a server.

The following tables provide recommendations for software and hardware configuration for the Web Storage System.

Table 15.6 Sample Web Storage System configuration

Storage Group	Database
First storage group	User mailboxes
First storage group	Executive mailboxes
Second storage group	Network news
Third storage group	Public folders

Table 15.7 Sample Web Storage System hardware configuration

Volume Label	Disk Controller	Data Bus	Disk Format
E	A	1	NTFS
F	A	1	NTFS
G	A	1	NTFS
H	B	2	NTFS
I	B	2	NTFS
J	B	2	NTFS

It is recommended that the C volume use RAID-1 technology and that it be reserved for the operating system and Exchange 2000 binary files.

Location for Mailbox and Public Folder Stores

You must allocate disk space in such a way that Exchange can quickly access the data it needs. In Exchange 2000, you can control the location of each database in a storage group and the placement of the transaction logs for the storage group by changing the properties of the storage group in System Manager.

For ease of management and maximum reliability, place databases and transaction log files on separate volumes, preferably with each database volume comprised of several disks. For maximum performance, place all databases in separate directories on one volume that is comprised of several disks. This configuration allows the system to balance the load evenly across the available disks.

For optimal performance, you should place the volumes on different disk controllers and different data buses. To provide fault tolerance and manageable disaster recovery, you should always place the transaction log files on a different volume from the corresponding databases. Placing the transaction log files on a separate and dedicated volume also increases performance; that way, because the transaction log files are sequential, the disk heads will always be in the correct position to write the next transaction.

Some examples of recommended configurations are shown in Table 15.8.

Table 15.8 Sample Web Storage System configuration

Storage Group	Database	Volume	RAID Technology	Disk Controller Cache Setting
First storage group	Mailbox stores	E	RAID-5	100% write
First storage group	Executive mailbox stores	I	RAID-5	100% write
First storage group	Transaction log files	F	RAID-1	100% write
Second storage group	Network news*	G	RAID-5	100% write
Second storage group	Transaction log files	H	RAID-1	100% write
Third storage group	Public folder stores	J	RAID-5	100% write
Third storage group	Transaction log files	H	RAID-1	100% write

*It is recommended that you use circular logging for network news data because the data is changing rapidly and recoverability is nut business critical.

RAID for Mailbox and Public Folder Servers

Mailbox stores and public folder stores are characterized by the random and balanced nature of disk read and write operations, whereas transaction log files are accessed sequentially, usually for write operations only.

The recommended RAID technology for mailbox stores and public folder stores is either RAID-5 (striped mirrors) or RAID-O+1 (mirrored stripes). Which technology you choose depends on a number of factors, which are listed in Table 15.9.

Table 15.9 RAID advantages and disadvantages

Technology	Advantages	Disadvantages
RAID-0+1	Fast read and write.	You must restore the entire volume if another disk fails during a rebuild.
RAID-5	Possibly less expensive than RAID-0+1. The difference depends on the number of drives that make up the volume.	You must enable write-back caching to provide adequate write performance and data protection. RAID-5 is slow for read and write. Rebuilding a RAID-5 volume takes longer than rebuilding a RAID-0+1 volume. You must restore the whole volume if another disk fails during a rebuild.

Although RAID-5 is often selected because it is less expensive, RAID-5 arrays are not always the most cost-effective solutions; RAID-5 can drastically reduce performance. Currently, 9-GB drives cost about the same as 18-GB drives, making RAID 0+1 a better choice for performance with a negligible difference in cost.

Note For drives that are 18 GB or larger, RAID-5 arrays are not recommended.

The recommended RAID technology for the transaction logs on mailbox arid public folder servers is RAID-1, because it offers the highest performance and sufficient reliability. Controller based RAID implementations are recommended, because they provide better performance and leave the processor to do other work. A configuration with the recommended RAID technologies is shown in Table 15.8.

Disk Controller Cache Settings for Mailbox and Public Folder Servers

To maximize RAID-1, RAID-0+1, and RAID-5 performance, the caches should be located on the disk controllers; disk drive caches should never be used because they are not protected and can result in database corruption. Due to the write-intensive nature of Exchange 2000, all the cache memory should be write-back cache. Read-ahead cache offers a performance gain when the disk

accesses contiguous blocks of information. Because disk operations in Exchange databases are non-sequential, it is unlikely that the disk controller can use the next disk block. Therefore, any transfer to read-ahead cache is wasted.

Read-ahead cache does not help when a stream of data (for example, a video clip) is read from the Web Storage System. The streaming file itself could be made contiguous on disk, but there is no guarantee that Exchange will save the video clip data sequentially inside the streaming file. The recommended disk controller cache settings appear in Table 15.8.

The settings in Table 15.8 are best for ordinary operation, but not for backup, when data is read from the volumes. However, by using multiple databases and by backing up the databases in a parallel fashion, you can reduce the backup time, even for large databases.

You configure the cache settings on the controllers; usually you cannot change them during operation. However some of the high-end controllers can adapt and increase the read-ahead cache, if this setting is enabled.

You should use the following methods to protect the disk controller caches:

- Error correcting cache (ECC) memory
- Mirroring of the cache (preferably on separate memory boards, so if you must change the controller, the cache data is still preserved on its own board)
- Swappable cache memory modules
- Battery backup

Also make sure that you are running the latest firmware recommended by the vendor. For more information about error correcting cache memory, see the Microsoft Product Support Services Knowledge Base articles Q151789, "XADM: Error -1018 (JET_errReadVerifyFailure)," and Q185075, "SAMPLE: PCIDMA.exe PCI Busmaster DMA Driver."

Paging File Settings for Mailbox and Public Folder Servers

Exchange 2000 Server is designed to make use of as much memory on the server as possible; this can put pressure on the Windows 2000 Cache Manager. Exchange 2000 is also designed to give up memory when other applications ask for it, but there are still times when memory paging is required. For this reason, follow these recommendations for paging files:

- Multiple paging files on multiple disks increase performance.
- Use the same settings for Initial Size and Max Size; this increases performance because it decreases the fragmentation of the paging file.

For more information about adjusting paging file size, see the *Microsoft Windows 2000 Sewer Resource* ***Kit.***

MIME Content

When you use Exchange 2000 in an Internet environment or when you save multimedia content in the Web Storage System, Exchange saves the data in streaming database files in Multipurpose Internet Mail Extensions (MIME) format. This enables high-performance access from Internet clients. However, the MIME format is an ASCII representation of the data; therefore it requires more disk space than native data or Rich Text Format (RTF) data. You should therefore plan for the additional disk space capacity when you move data to Exchange 2000.

Index Size

If you are using the content indexing feature, plan to reserve 25 to 30 percent extra disk space. Indexing large public folder stores can result in large indexes and high requirements for storage space. For better performance, keep the indexes on separate volumes. You can use System Manager to specify the location of the indexes when you enable content indexing on a mailbox store or public folder store.

Disk Space Margin of Safety

Volume capacity must accommodate future growth and potential problems. For example, the transaction log files volume should have enough capacity to hold two to three times the normal number of daily transaction log files. Extra disk space is necessary in case you encounter a massive increase in transactions because of message loops, virus attacks, or failure of the daily backup. Some utilities require as much disk space as the largest database because you need to save a copy of the database.

The single-instance message storage only occurs within a database. If you choose to implement several databases in one or more storage groups, you need more storage than for the equivalent number of users on Exchange 5.5.

Site Replication Service

If SRS is installed on the Exchange 2000 Server, it is recommended that you use a dedicated volume for both Srs.edb and Temp.edb. In addition, they should be in the same location as the working directory. Use RAID-5 and a 100-percent write disk controller cache.

To change the location for SRS

Caution Do not use a registry editor to edit the registry directly unless you have no alternative. The registry editor bypasses the standard safeguards provided by administrative tools. These safeguards prevent you from entering conflicting settings or settings that are likely to degrade performance or damage your system. Editing the registry directly can have serious, unexpected consequences that can prevent the system from starting and require that you reinstall Windows 2000 Server and Exchange 2000 Server. To configure or customize Windows 2000, use the programs in Control Panel or Microsoft Management Console (MMC) whenever possible.

1. From the **Start** menu, click **Run,** and then type **regedt32**.

   In the registry editor, select

   HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\MSExchangeSRS \Parameters.

2. Click **DSA Database file,** point to **Edit,** and then click **String.** Modify the volume and path for SRS.
3. Click **DSA Temporary file,** point to **Edit,** and then click **String.** Modify the volume and path for SRS.
4. Click **DSA Working Directory,** point to **Edit,** and then click **String.** Modify the volume and path for SRS.
5. Close the registry editor.

Disk Configuration for Connector Servers

A connector server is dedicated to running the following connectors, available with Exchange 2000:

- SMTP connector
- X.400 Connector
- Lotus Notes Connector
- Lotus cc:Mail Connector
- GroupWise Connector
- Microsoft Mail Connector

Typically, a connector server runs connectors of a single type, but if the traffic volume is low, you can run several types of connectors on the same server.

The architecture for the various connectors is shown in Figure 15.2.

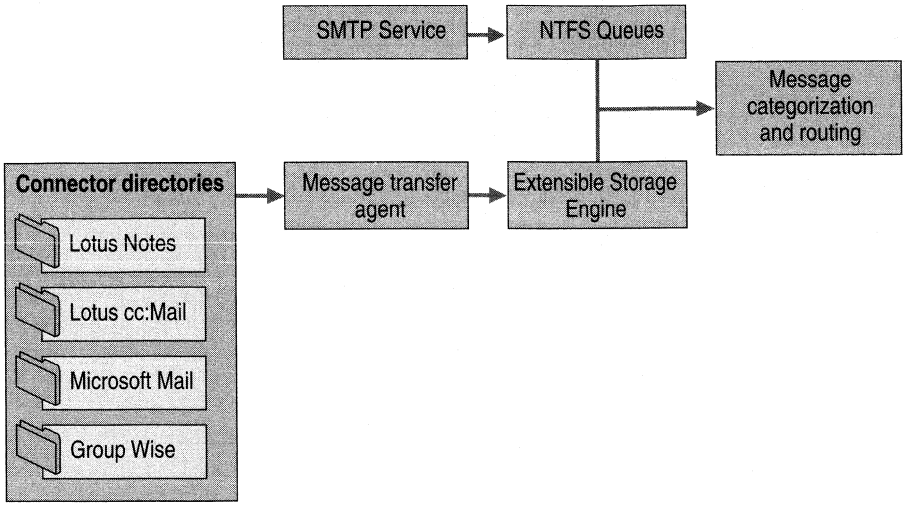


Figure 15.2 Overview of connector architecture

The illustration shows the connectors communicating with the MTA. Connectors, which run as Windows 2000 services, read and write to the corresponding directory. The MTA uses a separate directory for internal queues.

The SMTP service in Exchange 2000 provides connectivity to the Internet and acts as the default Exchange-to-Exchange server communication protocol. It communicates with the message categorization and routing components by using a dedicated directory.

The following section describes how to change the directory for different connector types.

> **Note** The SMTP service saves its configuration information in the IIS Metabase. The SMTP Service in Exchange 2000 is typically managed by using System Manager. The configuration information is written to Active Directory and then applied to the IIS Metabase by using the Exchange Metabase update mechanism. However, you cannot change the Queue directory by using the MMC snap-in, so you must update the Metabase directly. The Metabase can be configured through the Metabase Editor (MetaEdit) 2.0 utility found on the *Windows* 2000 *Sewer Resource* ***Kit*** companion CD.

To change the location for the SMTP service

1. From the **Start** menu, point to **Programs,** and then click **MetaEdit.**
2. Open LM, open **SmtpSvc**, open **1,** and then click **QueueDirectory**
3. Right-click **QueueDirectory**, and then click **Modify.**
4. In **Data,** type the new path.

When you change some of the settings, you must use Regedt32.exe to edit the Windows 2000 registry. Test these changes thoroughly before implementing them in a production environment.

> **Caution** Do not use a registry editor to edit the registry directly unless you have no alternative. The registry editor bypasses the standard safeguards provided by administrative tools. These safeguards prevent you from entering conflicting settings or settings that are likely to degrade performance or damage your system. Editing the registry directly can have serious, unexpected consequences that can prevent the system from starting and require that you reinstall Windows 2000 Server and Exchange 2000 Server. To configure or customize Windows 2000, use the programs in Control Panel or MMC whenever possible.

To change the location for the X.400 connector

1. In a registry editor, select HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\MSExchangeMTA\Parameters.
2. Click MTA **database path,** point to **Edit,** and then click **String.** Modify the path.
3. Close the registry editor.

To change the location for the Lotus Notes connector

1. In a registry editor, select HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\MSExchangeCoCo\Parameters.
2. Click **CONTROL_RootDir**, point to **Edit,** and then click **String.** Modify the path.
3. Close the registry editor.

To change the location for the Lotus cc:Mail connector

1. In a registry editor, select HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\MSExchangeCCMC\Parameters.
2. Click **Connector Store Path** entry, point to **Edit,** and then click **String.** Modify the path.
3. Close the registry editor.

To change the location for the Novell GroupWise connector

1. In a registry editor, select HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\MSExchangeCoCo\Parameters.
2. Click **CONTROL_RootDir**, point to **Edit,** and then click **String.** Modify the path.
3. Close the registry editor.

To change the location for the Microsoft Mail Connector

1. Change the *host*\maildat$ share on the connector server to point to the new directory.
2. Open **Windows Explorer** to the directory currently shared as **maildat$.** Remove sharing properties for this directory and create a directory in the new location. Share the new as **maildat$**.

Apart from the connector directories, the following additional components are important on a connector server:

- The message tracking log directory
- The Temp directory

Typically, message tracking is enabled on all components in an Exchange organization. Therefore, one or more entries are written to the message tracking log for each message that passes the server Because many components use the Temp directory when performing message format conversions, the Temp directory is used heavily on the connector server. It is also used extensively for full-text indexing. In addition, if the server runs content indexing, a public folder or mailbox server should have the Temp directory moved from the paging file disk.

To change the location for Message Tracking

1. In a registry editor, select HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\MSExchangeSA \Parameters.
2. Click **LogDirectory**, point to **Edit,** and then click **String.** Modify the path.
3. Close the registry editor.

Caution Do not use a registry editor to edit the registry directly unless you have no alternative. The registry editor bypasses the standard safeguards provided by administrative tools. These safeguards prevent you from entering conflicting settings or settings that are likely to degrade performance or damage your system. Editing the registry directly can have serious, unexpected consequences that can prevent the system from starting and require that you reinstall Windows 2000 Server and Exchange 2000 Server. To configure or customize Windows 2000, use the programs in Control Panel or MMC whenever possible.

To change the location of the Temp directory

1. Click **Start,** point to **Settings,** and then click **Control Panel.**
2. Open **System,** click the **Advanced** tab, and then click **Environment variables.**
3. In **System Variables** click **TEMP,** click **Edit,** and then type the new path.

As illustrated in Figure 15.2, the MTA communicates with the SMTP service through Extensible Storage Engine. Each message that passes between the MTA and the SMTP service generates one or more transactions in the corresponding transaction log file. In a scenario with high message volumes between these components, follow the recommendations from the earlier section, "Location for Mailbox and Public Folder Stores."

Location for Connector Servers

This section provides an example that illustrates the recommendations given thus far. In this scenario, the hardware is configured as shown in Table 15.10, along with the following connector types:

- SMTP service
- X.400 connector
- Lotus Notes connector

The connector server serves two X.400 connections and a large Exchange and Lotus Notes community. Therefore, the server experiences high traffic volumes on all three types of connectors.

Note The SMTP service is used for all Exchange server-to-server traffic.

The purpose of this is to place data in such a way that the different services (MTA, SMTP, and other connectors) can access as much data as possible at the same time.

For maximum performance, the directories should be placed on separate volumes. You can also place the directories on different disk controllers and different data buses.

An example of placement recommendations is shown in Table 15.10.

Table 15.10 Sample connector server configuration

Component	Directory	Volume	RAID Technology	Disk Controller Cache
SMTP Service	\EXCHSRVR\Mailroot	E	RAID-1	0% Read 100& Write
X.400 Connector	\EXCHSRVR\MTADATA	J*	RAID-1	0% Read 100& Write
Lotus Notes Connector	\EXCHSRVR\conndata	G	RAID-1	0% Read 10070 Write
Message Tracking logs	\EXCHSRVR\tracking.log	H	RAID-1	0% Read 100% Write
TEMP	\TEMP	I	RAID-1	50% Read 50% Write
Mailbox store and transaction log files	\EXCHSRVR\mdbdata	F	RAID-1	0% Read 100% Write

*Typically the SMTP service and the MTA will be the two most used components. For this reason it is recommended that you place these directories on different data buses.

The earlier section, "Location for Mailbox and Public Folder Stores," says that the transaction log files should always be placed on a different volume from the database for recovery reasons. However, the connector server is used only as a temporary interface mechanism, so the need for recovering the databases is not the same. Therefore, it is recommended that you do not use additional volumes to separate the transaction log files from the databases.

RAID for Connector Servers

The connector directories are characterized by high write performance. Almost all connector disk operations are write-only. Because read-only operations are random, the read cache is not a benefit.

The recommended RAID technology for the connector and other directories is RAID-1, because it offers the best performance and sufficient reliability. However, RAID-1 improves only read performance and does not significantly affect connector disks. Given the limited use of the mailbox store, RAID-5 is not necessary. A sample connector server configuration with the recommended RAID technologies is shown in Table 15.10.

> **Note** On servers with high rates of mail traffic, you can create a RAID-0+1 array for connector disks.

It is recommended that you use controller-based rather than software-based RAID implementations, because controller-based RAID provides better performance and frees the processor to do other work.

Disk Controller Cache Settings for Connector Servers

As described earlier, using controller-based caches improves performance. On a connector server, the read and write operations to the connector directories are not balanced; there are significantly more write operations than read operations. Therefore, you should divide the cache memory to allocate more memory for a write cache than a read cache. The recommended disk controller cache settings are shown in Table 15.10.

Antivirus Solutions

Most customers implement antivirus solutions in the Exchange environment. The typical solution is to place an SMTP and MIME virus scanner between the Internet and the connector server, either as a separate computer or as a solution running on an Exchange server by using the Exchange 2000 transport events to scan the messages. Regardless of which solution is used, follow the general guidelines regarding placement, RAID technology, and controller-cache settings. As a starting point, see the connector server guidelines, but remember to include the antivirus queue directory as one of the directories, because it needs special placement.

Disk Configuration for Front-End Servers

Exchange 2000 Servers that are configured only as front-end servers do not have significant storage requirements.

Disk Configuration for Data and Videoconferencing Servers

Exchange 2000 Conferencing Server offers services that do not have a large disk space requirement. Conferencing services only need certain mailboxes defined: the conferencing calendar mailbox and resource mailboxes. The conferencing calendar mailbox saves all conferences scheduled on the Windows 2000 site. The resource mailboxes contain the calendar for the resource, which is typically a limited amount of data. However, because the meeting request is sent to the resource mailbox, it also saves all information in the meeting request, such as attached documents.

For a dedicated data and video conferencing server, follow the disk configuration guidelines for a mailbox server. A sample configuration is shown in Table 15.11.

Table 15.11 Sample disk configuration for conferencing services

Storage Group	Database	Volume	RAID Technology	Disk Controller Cache Setting
First storage group	Conferencing mailboxes	E	RAID-5	100% write
First storage group	Transaction log files	F	RAID-1	100% write

In a configuration where an Exchange mailbox server is also running conferencing services, you can place the conference resources and the conference calendar mailbox in a separate mailbox store. By applying this strategy, you can minimize downtime if you need to restore scheduled conferences or move them to another mailbox store.

Disk Configuration for Instant Messaging Servers

Instant Messaging servers store little information an Exchange database. The storage requirements are shown in Table 15.12. Note that the storage requirements assume a maximum URL size of 256 bytes.

Table 15.12 Instant Messaging storage requirements

Item	Storage Requirement
User	540 bytes
User monitoring another user's status	875 bytes
Client logged on to an Instant Messaging home server	875 bytes

For example, consider the situation where you have 10,000 users with an average of 20 users monitoring each user. The total byte usage is ((540 + (875 * 20) + 875) * 10,000) = 183.5 MB of disk space. As mentioned earlier, this is a relatively low storage requirement.

The database used by the Instant Messaging servers is called Msimnode.edb and is placed in C:\Exchsrvr\IMdata by default. It is possible to change the placement of the database by using System Manager.

The non-sequential disk access pattern for Instant Messaging mailbox store resembles the disk access pattern for a mailbox store. Therefore, the disk configuration recommendations for Instant Messaging servers are similar to a mailbox server and are shown in Table 15.13.

Table 15.13 Instant Messaging server disk configuration

Component	Database	Volume	RAID Technology	Disk Controller Cache Setting
Instant Messaging	Msimnodes.edb	E	RAID-5	100% write
Instant Messaging	Transaction log files	F	RAID-1	100% write

Disk Configuration for Chat Servers

Exchange 2000 Chat Service does not have any significant storage requirements.

Memory Configuration

The amount of memory specified for each Exchange 2000 server type is for a server with no additional applications.

Memory for Mailbox and Public Folder Servers

Mailbox and public folder servers serve many who are active at any given time. This places a high requirement for processing power on the server. However, this varies with the number of users served by the server and the usage profile for those users.

Table 15.14 lists the recommended memory requirements for a given user population.

Table 15.14 Recommended memory configuration for mailbox servers

Number of Mailboxes	Memory Required
Fewer than 500	256 MB
500 to 1,000	512 MB
1,000 to 2,500	1 GB
2,500 to 5,000	2 GB

Memory for Connector Servers

You can use connector servers for a variety of protocols; the amount of memory that a connector server requires depends on which protocols it handles and the load for each protocol. The recommended memory configuration for connector servers is shown in Table 15.15.

Table 15.15 Recommended memory configuration for connector servers

Number of CPUs	Memory Required
2	256 MB
4	512 MB

Memory for Front-End Servers

The recommended memory configuration for computers designated only as protocol front-end servers is shown in Table 15.16.

Note The recommendations listed in Table 15.16 are based on the number of back-end mailboxes that the front-end server is supporting.

Table 15.16 Recommended memory configuration for front-end sewers

Number of Back-End Mailboxes	Memory Required
Fewer than 1,000	128 MB
1,000 to 3,000	256 MB
3,000 to 5,000	384 MB

Memory for Data and Video Conferencing Servers

The recommended memory configuration for data and video conferencing servers is shown in Table 15.17.

Table 15.17 Recommended memory configurations for data and video conferencing sewers

Simultaneous Users	Memory Required
Fewer than 500	256 MB
500 to 1,000	512 MB

Memory for Instant Messaging Servers

The recommended memory configuration for Instant Messaging servers is shown in Table 15.18.

Table 15.18 Recommended memory configurations for Instant Messaging servers

Instant Messaging Users	Memory Required
Fewer than 5,000	256 MB
5,000 to 10,000	512 MB

Memory for Chat Servers

The recommended memory configuration for Chat servers is shown in Table 15.19.

Table 15.19 **Recommended memory configurations for Chat servers.**

Chat Users	Memory Required
Fewer than 10,000	256 MB
10,000 to 20,000	512 MB

Network Configuration

It is important to configure the servers with the appropriate network bandwidth. In most cases this means 100-Mbps network interface cards and dedicated connections from the server to the network hub. Depending on the fault-tolerance required, servers are often configured with two network interface cards, each connecting to different network hubs.

Network Configuration for Mailbox and Public Folder Servers

In configuring networks for mailbox and public folder servers in large enterprises, it is common to use automated client/server backup technology; during backup and restore operations, the Exchange servers act as clients to the backup server. In such a configuration, you should use dedicated network interface cards and network hubs for the backup traffic. Isolating the backup and restore traffic from the user traffic provides the best performance for the backup server and for the users.

Network Configuration for Connector Servers

The typical configuration of a connector server is to use a single network interface card running at 100 Mbps.

Network Configuration for Real-Time Collaboration Servers

The typical configuration of real-time collaboration servers is to use a single network interface card running at 100 Mbps. However, the servers are often configured with an additional dedicated network interface card for an Internet connection.

CHAPTER 16

Message Routing

William Riedell, Senior Consultant, Microsoft

With Microsoft Exchange 2000 Server, you can design a messaging topology that closely fits your company's present and future needs. Although you still apply the same basic design strategies as with Microsoft Exchange Server 5.5 or earlier, new features in Exchange 2000 require careful consideration. This chapter discusses these features and the coexistence considerations between Exchange 5.5 and Exchange 2000 from a routing perspective.

In This Chapter

Exchange 5.5 Routing Basics

Exchange 2000 Routing Basics

Coexistence

Topology Considerations

Exchange 5.5 Routing Basics

Message transfer agents (MTAs) for Exchange Server version 5.5 and earlier are based on the X.400 standard for addressing and transporting e-mail messages. The X.400 standard supports several types of transport mechanisms, including Ethernet, X.25, TCP/IP, and dial-up lines.

You can group servers with full-time, reliable connections into Exchange sites. All servers in a site transfer messages directly to each other using remote procedure calls (RPCs), which creates a mesh network topology. To connect sites to one another, you can use the Site Connector, which is RPC based; X.400 Connector; or in some cases, Internet Mail Connector, which is based on Simple Mail Transfer Protocol (SMTP). Sites are determined by available network bandwidth or administrative boundaries. When a client running an earlier version of Exchange sends a message, one of the following processes routes the message:

1. Exchange determines if the destination address resides on the same server as the source address. If it does, the Microsoft Web Storage System places the message in the correct mailbox. MTA is not involved unless Exchange 5.5 Service Pack 1 or later is installed and message journaling is configured. Message journaling is configured on a database to create a copy of every message received by users on that database and to send the copy to a folder for archival purposes.

2. If the destination address is on another server on the site, MTA transfers the message directly to the destination server's MTA. This process involves RPC communications.
3. If the destination address is on a remote site, MTA routes the message to the remote site through a connector. There are two possibilities: either the address is on an adjacent site and can be reached by a connector, or the address is outside the organization and can be reached by a connector on this site or on another site.

Exchange 2000 Routing Basics

Exchange 2000 message transfer is based on SMTP. An Exchange 2000 server still contains an X.400-based MTA, but it uses it only to establish an X.400 messaging route between two Exchange routing groups, or between an Exchange routing group and an external X.400 system. When you compare SMTP as a native server-to-server transfer method in Exchange 2000 to the RPC method in Exchange 5.5, you see that SMTP is the more robust protocol. SMTP is also more reliable than RPC when you have slow connections. In Exchange 2000, SMTP is enhanced to expand the basic delivery functions of the protocol without compromising compatibility with other messaging systems. In Exchange 5.5 and earlier, the site concept defined both the administrative and routing topologies for an organization. In Exchange 2000, the site concept no longer exists. Instead, you use administrative groups and routing groups to separate the logical and physical organizations of your company. You use administrative groups to organize the servers according to the administrative structure of your organization and you use routing groups to map out the physical network of connected servers.

You can use one or more types of connectors to connect routing groups. In Exchange 2000, the three types of connectors are Routing Group connector, SMTP connector, and X.400 connector. Generally the Routing Group connector is used because it is the simplest to configure. However, you can connect routing groups with more than one connector of the same or different type for fault tolerance purposes. For example, you can connect two routing groups with one SMTP connector and two Routing Group connectors. If one Routing Group connector is not available, a server in the routing group can use the link state table to find the other SMTP connector and Routing Group connectors that you configured. The server in the routing group can then reroute information through the other connector.

Administration Groups and Routing Groups

Exchange 2000 uses administrative groups and routing groups instead of Exchange sites. You create administrative groups to define the administrative topology for large companies with a large number of departments or divisions, Exchange servers, and administrators. An administrative group is a collection of Exchange objects that are grouped together to simplify management of permissions; administrative groups can be administered separately from routing groups. For example, if a single mesh network has two distinct administration groups, one for sales staff and one for accountants, you can create two administrative groups. The Sales administrative group contains servers and objects managed by the Sales administrative team, and the Accountants administrative group contains servers and objects managed by the Accountants administrative team.

Routing groups are a group of servers that are constantly communicating with each other over reliable LAN connectivity. Information from one server to another server in the same routing group flows directly and immediately using Exchange's native SMTP protocol, not through a connector or on a schedule. A separate administrative team can administer routing groups.

If Exchange 5.5 is running, an administrative group is something different. In a mixed Exchange 5.5 and Exchange 2000 environment, administrative groups are equivalent to Exchange 5.5 sites for compatibility reasons. An Exchange 2000 installation runs in one of two modes: mixed mode or native mode. By default, when an Exchange 2000 server is added to an Exchange 5.5 organization, the organization runs in mixed mode. This ensures the future interoperability of the organization with Exchange 5.5 servers. You can also manually change the system from mixed mode to native mode by configuring the properties for the organization in System Manager. After the organization is in native mode, you cannot change it back to mixed mode.

Connectors

Exchange 2000 includes three connectors that you can use to connect routing groups: the Routing Group connector, the Simple Mail Transfer Protocol (SMTP) connector, and the X.400 connector.

Routing Group Connector

The Exchange Routing Group connector is an SMTP-based connector that is similar to the Exchange 5.5 Site Connector. However, if the Routing Group connector is used to connect to an Exchange 5.5 server, it automatically routes the message to the Exchange 5.5 server's MTA and uses RPC to transfer messages. The Routing Group connector, like the Exchange 5.5 Site Connector, enables single or multiple servers in a routing group to act as routing servers. The Routing Group connector also enables you to select a number of specific servers in the source routing group as source bridgeheads for the connector, instead of selecting only one server. Bridgehead servers facilitate communication outside the routing group. Configuring multiple servers within a routing group as bridgehead servers provides fault tolerance if one of the servers is offline.

SMTP Connector

The SMTP connector is primarily used to connect to the Internet or any other SMTP mail system. The SMTP connector is similar to the Exchange 5.5 Internet Mail Service, but it has additional features. For example, you can specify smart hosts instead of Domain Name System (DNS) lookup, outgoing MAPI content conversion, and remote delivery options.

The SMTP connector can use a smart host or mail exchanger (MX) records in DNS to route messages to the next server. If you designate a smart host, Exchange transmits information only to the smart host instead of repeatedly contacting the domain until a connection is made. This can be helpful when you are sending messages over the Internet and the remote domain can be reached infrequently or only during certain times.

In Exchange, messages sent by Internet clients are stored in Multipurpose Internet Mail Extensions (MIME) format, and no message conversion takes place when the messages are read. Messages sent by MAPI clients can be converted from Microsoft Rich Text Format (RTF) to MIME when Internet clients request them. SMTP allows you to set options when you deliver messages outside your Exchange organization. First, SMTP sorts messages and queues them for delivery, which allows SMTP to optimize connections by delivering multiple messages in one session. Then SMTP determines whether the receiving server is ready to receive messages. If the receiving server is not ready, the message remains in the queue and delivery is attempted again at intervals that you specify up to a maximum number of attempts. All outgoing messages, if not sent to a connector, can be sent to a smart host, which can then send them to remote recipients. You can also configure a list of trusted external DNS servers that your virtual server uses when processing messages addressed to non-local recipients. The virtual server uses only these DNS servers to look up remote addresses.

X.400 Connector

You can use the X.400 connector to connect Exchange 2000 routing groups or to connect the Exchange 2000 system to an external X.400 system. Using an X.400 connector to link routing groups together is recommended only if X.400 connectivity exists and there is a severe shortage of network bandwidth between two locations. The message payload of a large attachment transferred over an X.400 connector is less than the message payload of the same attachment when it is base 64 encoded for SMTP transfer.

Routing functionality and transfer of link state information are the same for an X.400 connector as for a Routing Group connector; however, unlike a Routing Group connector, you can define only a single host for the local and remote bridgehead servers. This means that load balancing can be achieved only with multiple X.400 connectors. An X.400 messaging route defines the path that an X.400 message follows to reach its final destination. It allows the MTA in both routing groups or systems to communicate. You can configure multiple X.400 connectors to support different transport types, such as TCP/IP, remote access, or X.25.

Link State Routing

Exchange 5.5 relies on the Gateway Address Table (GWART) for route selection in the organization. This method is flawed because it uses a distance vector routing algorithm, which in certain cases, is susceptible to routing loops. Exchange 2000 solves this problem by using a link state routing algorithm and a protocol to propagate link state information around the Exchange 2000 organization. The algorithm and the protocol are collectively referred to as the link state algorithm. The link state algorithm notifies other servers in the organization, in real time, about the state of the network. Link state enables an Exchange 2000 server to select the most efficient path for a message based on the current state of the network. This provides the following two advantages:

- The most efficient routes are selected at the source with complete knowledge of the current state of the network.
- As an advanced algorithm, link state routing is guaranteed to be free of loops.

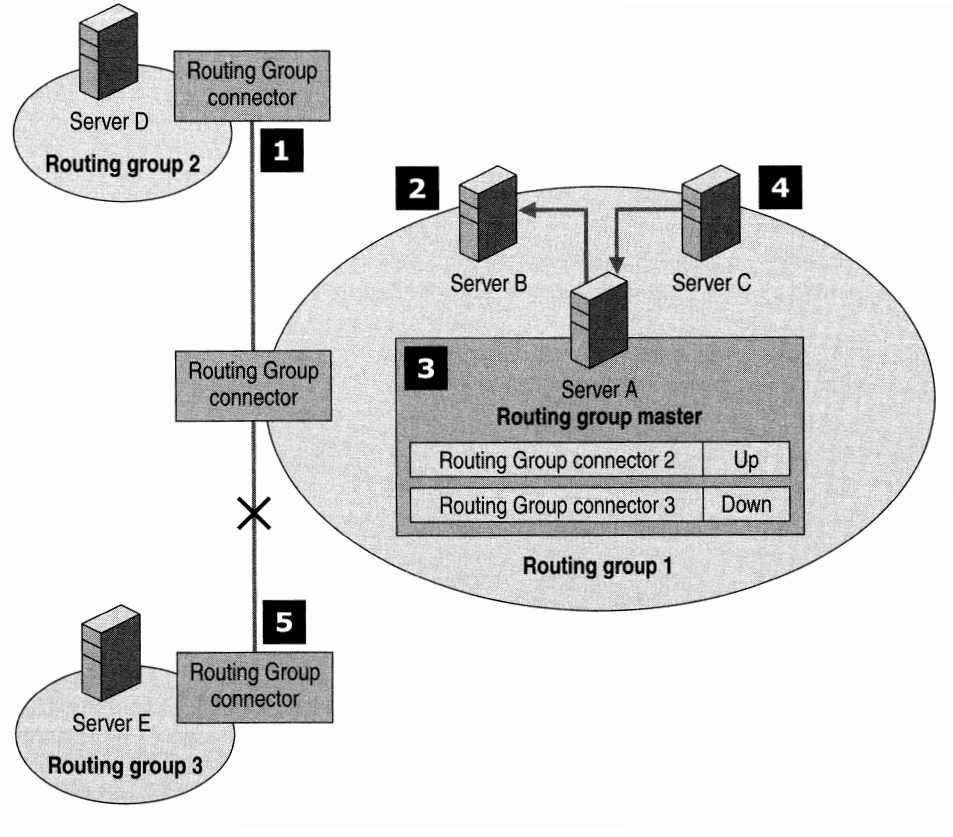


Figure 16.1 The link state algorithm is used to intelligently route messages

Figure 16.1 illustrates how the link state algorithm is used when a connection goes down. The Exchange 2000 connection servers that are actually at the connection ends are the first servers to notice that a link is unavailable. When a connection goes down, Exchange takes the following steps:

1. In this example, server C notices that the connection to routing group 3 is down. Between routing groups, connections are established using port 25 for SMTP.
2. Server C communicates the link state information to server A, which is the routing group master.
3. Communication in a routing group is established on port 691. The routing group master then updates the link state table.
4. The routing group master communicates the changes to all servers within its routing group.
5. All servers that have external connections in their routing group propagate that information to the entire organization through routing group propagation.

Once a link has been tagged as being unavailable or down, the connecting server continues to try to establish a link at the interval specified for normal message retry on the SMTP virtual server. Once a link is established, the link state table is updated to reflect that the link is now available. If a message cannot reach its destination within the expiration time-out duration specified on the SMTP virtual server, the sender receives a non-delivery report (NDR).

If for some reason the routing group master server is unavailable, the remaining servers in the routing group continue to operate with the information they had before they lost connectivity to the master server. When the routing group master server comes back online, it reconstructs the link state table by communicating with other servers in the routing group. As the link state table is updated, the information is communicated to the entire routing group. If the master server will be offline for a long period of time, you can manually configure another server in the routing group to be the master server. By default, the routing group master server is the first server in a routing group.

Route Selection

When a client sends a message, Exchange performs the following tasks to correctly route the message:

- If the destination address resides on the same server as the source address, the Exchange server delivers the message to the recipient's mailbox.
- If the destination address is on another server in the routing group, the message is handled either by the SMTP service (if the destination server is Exchange 2000) or by the MTA. If the destination server is an Exchange 5.5 server, the MTA transfers the message directly through RPC calls.
- If the destination address belongs in a remote routing group, which is determined through a lookup in Active Directory, Exchange identifies the route to take and routes the message to the appropriate connector.
    - Before it selects a route, the Exchange 2000 server examines its copy of the link state table to determine if any connections are down and which connections are up.
    - If the destination is located outside the Exchange organization, the message is routed to the Exchange server that hosts the outbound connection. If that connection is down, Exchange tries to reroute the message to another bridgehead server.
- This process is repeated each time the message is handled by an SMTP service or the MTA.

Coexistence

Most organizations migrate from earlier versions of Exchange and during the migration Exchange 5.5 and Exchange 2000 coexist in a mixed environment. This section discusses routing in a mixed Exchange site or routing group and how to route messages between sites and routing groups.

You can install an Exchange 2000 server on an existing Exchange 5.5 site, but you must have at least one computer running Exchange 5.5 Service Pack 3 in the site. When an Exchange 2000 server is installed on an Exchange 5.5 site, one administrative group and one routing group are created in Active Directory; these groups share the name of the original site. In most cases, the Exchange 5.5 server that was previously called Routing Calculation Server (also called Routing Information Daemon [RID]) generates a Gateway Address Resolution Table (GWART) to be used by the Exchange 5.5 servers to propagate connector information. If there are no Exchange 5.5 servers in an administrative group or site, the Exchange 2000 routing group master can generate the GWART in Active Directory, and then Site Replication Service replicates the GWART to the Exchange 5.5 directory in adjacent sites. All Exchange 2000 group and connector information is stored in Active Directory and is replicated to the entire organization. Exchange 2000 does not use GWART for routing. Exchange 5.5 uses it to inform the Exchange 5.5 servers that they can route messages to Exchange 2000 servers and connectors on the same site, and vice versa.

If there are multiple Exchange 2000 servers on an Exchange 5.5 site, they communicate with one another through the SMTP protocol. An Exchange 2000 server, on the other hand, transfers messages to an Exchange 5.5 server through RPC. This happens if an organization uses earlier connectors that have not been updated to run on Exchange 2000 or if you need Professional Office System (PROFS) or SNA Distribution System (SNADS) connectivity, which is not provided in Exchange 2000.

Within a mixed Exchange site, you can upgrade your bridgehead server to Exchange 2000 even if you are connecting Exchange 5.5 sites. After your hub or multiple bridgehead servers are upgraded to Exchange 2000, they begin to communicate by using SMTP. This increases throughput and provides resilience to slow network connections.

Topology Considerations

Whether coexisting and migrating from an earlier Exchange 5.5 organization or building a new organization, you should take the following design considerations into account:

- In native mode, Exchange 2000 can be extremely dynamic because you can change the memberships of the routing groups and administrative groups. This is useful when administration models or network topologies change.
- Exchange 2000 provides load balancing in the form of a round-robin DNS between servers, both sources and targets. A round-robin DNS is a mechanism that directs incoming requests to servers on a rotating basis. This is done by looping through a list of IP addresses belonging to the servers in the configuration. When an e-mail client attempts to access a mailbox on an Exchange server, the client is given the first IP address on the list. The second client request is given the second IP address in the list, and so on. If there are four servers on the round-robin list, all four IP addresses are used before the first IP address is used again, and the loop starts over. In addition, Exchange 2000 offers Improvements over the Exchange 5.5 Site Connector if one of the source bridgehead servers is down—Exchange connectors automatically try not to use that server until it comes back up. If there are multiple connectors with the same cost, each server picks a random connector and uses it for a period of time. Over multiple servers, this functionality simulates round-robin behavior.
- When a client must use an alternate server to access public folder content, Exchange 2000 uses routing groups to calculate the closest available server, which it determines by using a cost property set on the Routing Group connector. The cost for each Routing Group connector is stored in a single cost database that is shared with e-mail routing calculations. Redundant cost tables maintained in previous versions of Exchange are eliminated. In Exchange 5.5, site affinities were not transitive. For example, if you establish an affinity between site 1 and site 2, and between site 2 and site 3, you do not automatically get affinity between site 1 and site 3. In Exchange 2000, affinities are transitive. If all routing groups in an organization are connected to allow e-mail to flow, all servers receive public folder referrals. If a server does not contain replicas of public folders, you can mark specific Routing Group connectors to deny public folder referrals; the client contacts the next server in the referral list.

- Because SMTP is used in a routing group and SMTP is more resilient over slow links, routing groups can span slow networks more effectively than RPC, which is the native transfer mechanism for Exchange 5.5. Thus, your organization may be able to deploy fewer routing groups by using Exchange 2000 than by using sites with Exchange 5.5.
- SMTP mail is not compressed; though you reduce network bandwidth requirements with SMTP, the increased load on the Exchange server CPU would reduce server performance.
- Unlike Exchange 5.5, where intra-site and inter-site RPC communications are encrypted, SMTP communications in Exchange 2000 are not encrypted. However, each server uses SMTP authentication with Kerberos. There are two options for encryption, although neither is done by default: Internet Protocol security (IPSec), which is built into Microsoft Windows 2000, and Transport Layer Security (TLS) encryption, which is built into the SMTP service.
- Upgrading an Exchange 5.5 hub, in a hub-and-spoke design, can significantly increase hub performance and stability because performance is enhanced by Exchange 2000 and the use of SMTP.
- Link state routing is most effective in an organization with multiple routing groups using multiple paths. In a traditional hub-and-spoke design, you do not see as much improvement as you do in a mesh network or a network with multiple connections.
- Adding routing groups increases the size of the link state table and increases potential link state status messages. This can increase the amount of link state message traffic.

Exchange no longer includes a Dynamic RAS connector. Instead, Exchange 2000 uses the on-demand capabilities in Windows 2000, specifically, the combination of the Routing Group connector and the Routing and Remote Access demand-dial interface.

CHAPTER 17

Backbone Configuration and Tuning

Cherif Djerboua, Consultant, Microsoft
Jens Trier Kasmussen, Principle Consultant, Microsoft

The underlying transport and routing engine of a messaging system is critical to the system's success. This core component is fundamental to the operation of an efficient and reliable enterprise messaging system, especially if users are distributed around the globe.

By using Simple Mail Transfer Protocol (SMTP) as the native communications protocol between servers running Microsoft Exchange 2000 Server, many new opportunities emerge; at the same time, some of the deployment issues with earlier transport implementations disappear. For example, companies with a distributed user base normally design their Exchange site model around the availability of network bandwidth rather than the desired administration model. Because Exchange 2000 no longer uses remote procedure call (RPC) for message transfer, you can now devise a more flexible routing scheme. In addition, Exchange 2000 is organized into administration groups and routing groups instead of sites. This gives you more flexibility in both administration and routing.

In This Chapter

Directory Access

Routing and Transport

Public Folders

Address Lists

Directory Access

Exchange 2000 uses Lightweight Directory Access Protocol (LDAP) to query and update Active Directory directory service. Because this is a standard protocol, you can view traffic between an Exchange 2000 server and the global catalog server by using Network Monitor (unless Exchange 2000 is installed on the global catalog server). In addition, Directory Service Access (DS Access) uses the Exchange System account to authenticate users, to obtain full rights to create directory entries, and to read and write from Active Directory. However, the LDAP data itself is not encrypted, so if you look at the traffic on the network you can see all the Exchange read and write activity with Active Directory.

Exchange 2000 System Manager uses Active Directory as a storage mechanism for Exchange configuration information.

Clients can access Active Directory information by communicating directly with a global catalog server or by sending a request through Exchange 2000. Active Directory supports both LDAP and MAPI queries (for backward compatibility with older MAPI clients) and Exchange 2000 provides a request process for Microsoft Outlook Web Access clients who make directory queries by using HTTP.

Directory Service Access API

All Exchange 2000 services that require access to Active Directory (whether for reading configuration information or writing new entries to the directory) use the Directory Service Access (DS Access) application programming interface (API). The most important part of DS Access is the shared cache, which caches search results between different services in Exchange.

Specifying an Active Directory Server

The DS Access process needs to contact Active Directory servers for two primary reasons: to look up information in the address book and to read configuration data.

To contact an Active Directory server, the DS Access process makes a Domain Name System (DNS) query. For global catalog access, DS Access queries the Windows 2000 site to which the Exchange 2000 server belongs. If all global catalog servers in the site are unavailable, DS Access queries other sites. For domain controller access, DS Access first queries domain controllers within the same site and domain. If no domain controller is available, DS Access queries outside the site but within the same domain to find a domain controller. If more than one domain controller is available, DS Access uses the round-robin method to choose one.

If the required results are not stored in one of the domain controllers, the DS Access process makes a DNS query for the nearest global catalog server and requests the information.

You can force an Exchange 2000 server to always use the same Active Directory server; if the Active Directory server does not exist on the network, it reverts to DNS. This is done by changing the following registry settings.

Caution Do not use a registry editor to edit the registry directly unless you have no alternative. The registry editor bypasses the standard safeguards provided by administrative tools. These safeguards prevent you from entering conflicting settings or settings that are likely to degrade performance or damage your system. Editing the registry directly can have serious, unexpected consequences that can prevent the system from starting and require that you reinstall Exchange 2000. To configure or customize Exchange 2000, use the programs in Control Panel or Microsoft Management Console (MMC) whenever possible.

To specify a server on which to look up configuration data

1. On the **Run** line, type **regedt32.exe** or **regedit.exe**, and then click OK.

   Navigate to HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\MSExchangeDSAccess\Instance0.

   Note Create the **Instance0** subkey if it does not exist.

2. Select or create the **ConfigDCHostName** entry.
3. To specify a single server for all configuration data look ups, change the REG–SZ value to *servername.domainname.com.*

   In Regedit.exe, right-click the entry, and then click **Modify.**

   –or–

   In Regedt32.exe, click the entry, click **Edit,** and then click the appropriate menu choice.

4. Select or create the **ConfigDCPortNumber** entry.
5. Change the REG–DWORD value to **0x389** (the LDAP port number).
6. Close the registry editor.

To specify a sewer on which to look up addresses

1. In a registry editor, navigate to HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\MSExchangeDSAccess\Profiles\Default.

   Note Create the subkey if it does not exist.

2. Select or create the **UserGC1** entry.
3. To specify a single server on which to look up addresses, change the REG–SZ value to sewername.domainname.com.
4. Select or create the **PortNumber** entry.
5. Change the REG–DWORD value to **0x3268** (the global catalog server port number).
6. Select the **IsGC** entry.
7. Change the REG–DWORD value to **0x1**.
8. Close the registry editor.

Directory Cache

Part of the DS Access API is the directory cache. The cache is enabled by default and memory is dedicated to the cache when you start Exchange 2000. All Exchange system queries made through the DS Access API are cached for a total of 10 minutes or until the cache memory has reached 4 MB (each cache entry takes approximately 3.5 KB of space).

By caching Active Directory queries, the Exchange 2000 server improves responsiveness and decreases the load on the domain controllers.

Cache Parameters

The DS Access API cache can be configured by using the registry editor.

> **Caution** Do not use a registry editor to edit the registry directly unless you have no other alternative. The registry editor bypasses the standard safeguards provided by administrative tools. These safeguards prevent you from entering conflicting settings or settings that are likely to degrade performance or damage your system. Editing the registry directly can have serious, unexpected consequences that can prevent the system from starting and require that you reinstall Exchange 2000. To configure or customize Exchange 2000, use the programs in Control Panel or Microsoft Management Console (MMC) whenever possible.

To change settings for the DS Access cache

1. In a registry editor, navigate to HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\MSExchangeDSAccess.
2. Select or create the **CachingEnabled** entry.
3. To enable caching for DS Access, assign a value of **0x1**.

   To disable caching for DS Access, assign a value of **0x2.**
4. In a registry editor, navigate to HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\MSExchangeDSAccess\Instance0.

   > **Note** Create the **Instance0** subkey if it does not exist.
5. Select or create the **CacheTTL** entry.
6. To set the amount of time that items exist in the DS Access cache, assign a value (0x600 represents 600 seconds, the default value).
7. Select or create the **MaxMemory** entry.

8. To set the amount of memory available for items in the DS Access cache, assign a value (0x0 represents 4096 KR, the default value).
9. Select or create the **MaxEntries** entry.
10. To set the maximum number of entries in the DS Access cache, assign a value (0x0 allows an unlimited number, the default value).
11. Close the registry editor.

At present, there is no tool for viewing entries in the cache. To troubleshoot problems with the cache, the first step is to disable the cache arid determine if this resolves the problem.

Routing and Transport

This section discusses concepts and strategies for optimizing the routing and transport functions in Exchange 2000.

Advanced QueuingEngine

With the advanced queuing engine, Exchange 2000 extends the base SMTP protocol that is installed with Windows 2000 Server. The advanced queuing engine routes messages through the system to their proper destination. It maintains queues for inbound and outbound messages, calls transport event sinks, and generates delivery status notifications for acknowledgement or error in message delivery.

Message Categorizer

The message categorizer resolves sender and recipient addresses in Active Directory; this includes looking at certain attributes of senders and recipients and expanding distribution lists. The messages might be destined for a local mailbox store, a remote host available through the message transfer agent (MTA), or a remote host available over SMTP. The message categorizer is a component of the advanced queuing engine, and its behavior can be modified through the Categorizer Transport Events.

The message categorizer looks up recipients in Active Directory, determines the server on which the recipient's mailbox exists, and routes the mail to that server. If the server matches the local server's name, the advanced queuing engine sends the message directly to the mailbox store.

Distribution Group Expansion

Distribution lists do not exist in Exchange 2000. They are called distribution groups and are expanded by the message categorizer in the following ways:

- The message categorizer expands distribution groups locally.
- The message categorizer detects loops in distribution lists.
- The message categorizer can force the mail-enabled groups to be expanded on a specific server.

A given distribution group might need to be expanded on a specific host. This host attribute exists for each distribution group. If the distribution group expansion host attribute is blank, any host can expand that distribution group. The attribute contains a globally unique identifier (GUID) for a specific SMTP virtual host that expands the mail-enabled group. The attribute is called Home-VSI (for Home-Virtual Server Instance). The directory service in Exchange Server version 5.5 and earlier uses msExchExpansionServerName and the Home-MTA property. Both of these properties contain the domain name of an MTA that expands a distribution list. When the Home-VSI is replicated to a directory service in Exchange 5.5 or earlier, Home-VSI is turned into a Home-MTA or msExchExpansionServerName and vice versa. Home-VSI is set only on distribution groups.

Connector Types

There are two major connector types:

- Routing Group connectors and SMTP connectors, which bypass the MTA and use SMTP directly. Routing Group connectors look at the version of the target server. If the target server is running Exchange 2000, the connector bypasses the MTA and uses SMTP. If the target server is running Exchange 5.5 or earlier, the routing group connector routes the message to the MTA. The MTA treats the connector as an Exchange 5.5 Site Connector and sends messages by using RPC.
- X.400 connector and EDK gateway, which use the MTA as an integral part of the message path.

Table 17.1 lists the different types of connectors and whether they are available in Exchange 5.5 and/or Exchange 2000.

Table 17.1 Exchange connectors

Name of Connector	Present in Exchange 5.5	Present in Exchange 2000
Site Connector	Yes, as an extension of the MTA	No, replaced by the Routing Group connector, which uses RPC or SMTP, depending on whether it connects to Exchange 5.5 or Exchange 2000, respectively
X.400 connector	Yes, as an extension of the MTA	Yes, as an extension of the MTA
Internet Mail Service	Yes, interfaces with the MTA	No, replaced by the SMTP connector
Lotus cc:Mail and Microsoft Mail	Yes	Yes
SNA Distribution System (SNADS) and Professional Office System (PROFS) connectors	Yes	No
Dynamic Remote Access Server connector (DRAS)	Yes	No, replaced by the dial-on-demand X.400, Routing Group, or SMTP connectors

For a complete discussion about connectors, see "External Connectivity" in this book.

Routing Group Connector

The Routing Group connector is analogous to the Site Connector in Exchange 5.5. It is easy to configure and it establishes mail connectivity between two routing groups. When you define multiple routing groups, you need to deploy a connector to enable messaging; the Routing Group connector is the easiest connector to configure.

The Routing Group connector, like the Site Connector, uses a protocol that is not explicitly defined. That is, you don't need to worry about what protocol it is using, because there aren't any protocol-level settings on the connector object to the Routing Group connector, which uses SMTP.

A Routing Group connector can also connect an Exchange 2000 server with an Exchange 5.5 server. In this case, it looks like a Routing Group connector on the Exchange 2000 side, but the Exchange 2000 server automatically determines that the other server is running an earlier version of Exchange and sends the message using the MTA and RPC. To the Exchange 5.5 server, the connection appears to be from another Exchange 5.5 server.

SMTP Connector

The SMTP connector is analogous to the Internet Mail Service in Exchange 5.5. It has two primary purposes:

- To connect an Exchange 2000 organization or routing group (depending on the scope) to the Internet or to an external mail system. In this case, you simply add an address space for the Internet addresses that might travel over the connector. In most cases, it will be an asterisk (*) with an address type of SMTP.
- To connect two routing groups when you want to specify that SMTP be used. The major difference between an SMTP connector and a Routing Group connector is that the Routing Group connector uses SMTP between Exchange 2000 servers and it uses RPC between an Exchange 2000 server and an Exchange 5.5 server; the SMTP connector always uses SMTP. If the server at one end of the SMTP connector is an Exchange 5.5 server, it must have Internet Mail Service configured. In addition. the SMTP connector gives you special control over SMTP protocol-level parameters, such as SMTP authentication or encryption. To use an SMTP connector to connect routing groups together, instead of an address space, you must use the **Connected Routing Groups** tab to indicate the routing group on the other end of the connector. You can combine both address spaces and connected routing groups on one connector.

The SMTP connector also supports two commands relating to triggered message delivery (ETRN and TURN).

X.400 Connector

The X.400 connector is provided for the following reasons:

- Connectivity to other external X.400 MTAs
- Connectivity to an X.400 service provider
- Connectivity between two routing groups

The X.400 connector in Exchange 2000 has changed little since previous versions of Exchange. There are some architectural changes, including the switch to LDAP directory lookups.

> **Note** X.400 connectors over Transport Class 4 (TP4) protocol are no longer supported because Windows 2000 does not provide network support for TP4.

You can use the X.400 connector to link routing groups together. This is desirable where only X.400 connectivity exists, or where the network link is unreliable or has limited network bandwidth. The size of a large message attachment transferred over an X.400 connector is smaller than the same attachment that has been Base64 encoded for SMTP transfer. In addition, the X.400 MTA allows for graceful recovery of associations when transient problems exist on the network.

When you use an X.400 connector between routing groups, MTAs pass link state information by sending a binary object before messages are transferred. However, unlike a Routing Group connector, only a single host can be defined for the local and remote bridgehead servers. This means that load balancing and bridgehead server fault tolerance can be achieved only by configuring multiple X.400 connectors.

Routing and Link State Information

The Exchange 2000 Server MTA does not perform routing functions, but instead a calls upon the Exchange 2000 routing engine for routing decisions. The routing engine uses link state tables for its routing decisions. These tables are built using the link state algorithm, which has been used extensively on the Internet for many years in the form of Open Shortest Path First routers, which reside at the network layer. However, the link state table in Exchange 2000 works at a much higher level in the Open Systems Interconnection (OSI) model; thus, it does not rely on the underlying network infrastructure and protocols.

The link state algorithm propagates the state of the messaging system in almost real time to all servers in the organization. Exchange 2000 uses routes and costs just as Exchange 5.5 does, but the link state algorithms provide the following advantages:

- Each Exchange server can make the best routing decision at the source instead of sending messages down a path for which a link is unavailable.
- Messages no longer bounce back and forth between servers because each Exchange 2000 server has information about whether alternate or redundant links are available.
- Message-looping problems no longer occur.

The MTA in Exchange 2000 is modified to support Active Directory by using LDAP calls.

How Link State Works

Link state information is most effective when working with multiple routing groups in an organization, especially where redundant paths are available. Each routing group has a master server assigned by the administrator that receives link state information from different routing groups. The master keeps track of this data and propagates it to the rest of the servers in the routing group. The master ensures that all servers in the routing group advertise the same information. The master is normally the first server to be installed in the routing group, but you can change this in Exchange System Manager.

There are differences between the ways that link state data is propagated between routing groups and within a given routing group. Between groups, new information is relayed through SMTP on port 25 and is sent between servers when a change takes place. Within a routing group, the routing master sends and receives link state information on TCP port 691. When a non-master receives new link state information, it immediately transfers it to the routing master, so other servers are alerted to the routing change. There are only two states for any given link: up or down. Connection information, such as whether a link is actively sending or in a retry state, is not propagated; it is known only by the server involved in the message transfer.

All link state data received is held in the memory of the Exchange 2000 servers and not written to disk. If any one of the servers restarts (including the routing master), the server receives the information from other servers.

Link State Updates

Within a routing group, each server communicates with the master using a TCP-based link state table protocol. Each server (including the master) monitors TCP port 691 (registered with Internet Assigned Number Authority [IANA] for this purpose). The master broadcasts changes only to the servers in its routing group.

When servers in two routing groups communicate using SMTP, they use a version of the link state algorithm that works as an extension to SMTP. Exchange 2000 servers advertise the XLINK2STATE protocol; when one Exchange 2000 server sees another advertising that routing information has changed, it attempts to trade routing information. Routing information transfers only if the two servers are in the same organization and if routing information differs between servers.

When servers in two routing groups cominunicate over X.400, they also use link state information. The MTA constructs a special X.400 message to transfer this information.

When a bridgehead server sees that a link is unavailable, it tags that connector as down and sends the data to the routing group master. The master immediately informs the other servers in the routing group of this change, and it also propagates this information to other routing groups. The sample that follows shows a sample network monitor trace of a link state update between a member and a master.

Trace 17.1 Member - master link state update

```
00000030                    7B 30 30 30 30 30 30 35 31 7D
{00000051}
00000040  20 56 53 5F 43 4F 4E 4E 20 37 66 65 31 32 65 66
.VS_CONN.7fe12ef
00000050  65 65 38 66 31 36 35 34 64 62 63 32 37 31 37 35
ee8f1654dbc27175
00000060  38 35 35 64 37 66 37 64 31 20 63 33 38 30 61 62
855d7f7d1.c380ab
00000070  30 66 37 62 37 62 66 66 34 33 38 65 66 62 38 37
0f7b7bff438efb87
00000080  30 37 35 36 38 63 66 38 62 39 20 44 4F 57 4E 20
07568cf8b9.DOWN.
00000090  20
```

The information contained in this trace is interpreted as follows:

- The number, 7fe12, is the GUID of the connector affected.
- The third hexadecimal number (c380a) is the GUID of the virtual server that noticed the change.
- The final word (DOWN) sets the status of the connector.

As noted above. between routing groups, data is transferred over SMTP on port 25. The format of the data is roughly similar to that of the intra-routing group communications. If you perform a network trace for link state data, you will notice the X-LINK2STATE command verb is used to denote the type of data, and the information is sent in chunks.

The *Exchange* 2000 *Resource Kit* companion CD provides a tool called WinRoute that connects to the link state port (port 691) on an Exchange 2000 server and extracts the link state information for the Exchange organization. Although the information is usually a series of GUIDs, WinRoute matches the GUIDs of connectors and bridgehead servers to objects in Active Directory and shows you an easy-to-read version of the information. The upper half of the window shows the interpreted data, the lower half shows the raw data from the link state port.

Message Delivery to Remote Servers

The Exchange 5.5 MTA communicates directly with other servers in the site, across a Site Connector or an X.400 connector. The Exchange 2000 MTA communicates directly only with other X.400 MTAs and Exchange 5.5 servers.

The MTA communicates with another MTA in the routing group or across a connector in order to pass mail messages. After certain parameters such as Window Size, Checkpoint Size, and Recovery Timeout are agreed upon between the MTAs, mail messages flow.

Exchange 2000 servers do not use a Gateway Address Resolution Table (GWART) file for routing purposes. However, a GWART table is generated and replicated by Active Directory Connector (ADC) if an Exchange 2000 server is designated as the routing calculation server or if there is no Exchange 5.5 server in an Exchange site.

Routing Group Scenario

The following scenario discusses routing group architecture, the choice of connectivity, and connection agreements.

Figure 17.1 illustrates the planned Windows 2000 site and Exchange routing group topology and the network links deployed between the locations of a fictitious company.

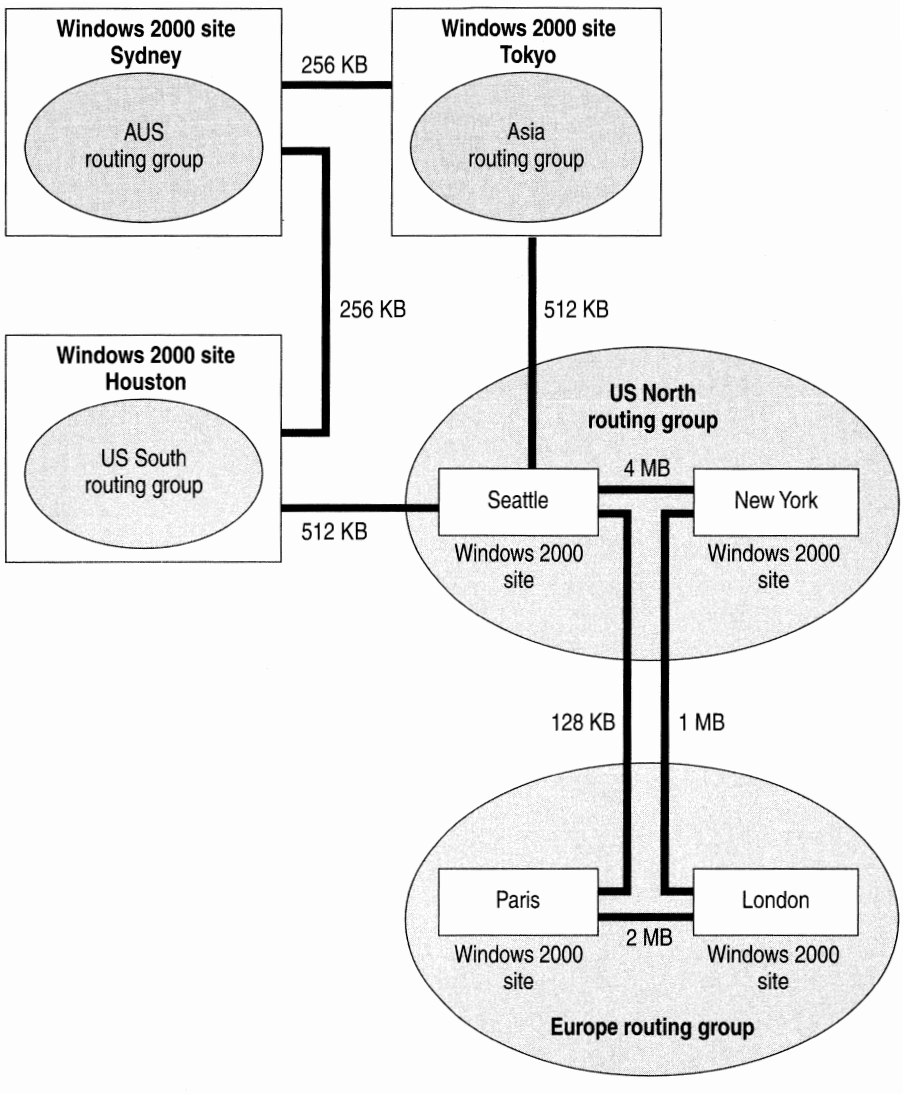


Figure 17.1 Planned topology and network connectivity between locations

Even when there appears to be a high-bandwidth link between two locations, you must analyze other traffic that may be present on the line to calculate the net available bandwidth. In addition, you need to ask the following questions:

- Are there certain groups of users that send large messages to one another on a regular basis?
- What's the average size of messages across the network'?
- Which public folders do users access?

Users have access to all public folders within a routing group. This means that if a routing group spans a slow WAN connection, client public folder access could slow to an unacceptable level.

Many of the design philosophies used to design the site architecture in Exchange 5.5 and earlier can be used to design Exchange 2000 routing groups. After you perform the network analysis and clarify the business requirements, the next step is to lay out the boundaries. You can group multiple locations together to form single routing groups, although this could affect client connections to public folders. In Figure 17.1, although USA North and USA South could be part of the same routing group, thit could result in slow public folder access.

Routing Group connectors are used in this scenario because they offer the best functionality and resilience. Each Routing Group connector hat source and target bridgehead servers. To get the maximum efficiency from the network, configure the bridgehead servers so all connections occur over a single network link. For example, in Figure 17.1 the target bridgehead servers specified in the connector from USA South to USA North include only connector servers in Seattle, because there is no direct network connectivity between Houston and New York.

When message transfer must take place, a local bridgehead server reads the list of target servers and chooses the one with the lowest connection cost, or if the servers have the same cost, a random one. If the first server is down, the local bridgehead server uses another random server. and so on. Subsequent messages use the same algorithm.

Different messaging costs can also be implemented to dictate the primary messaging path. In the scenario, the Asia and AUS routing groups have a higher cost link between them, so messages between London and Sydney travel over the lower cost network link through Houston.

Low-Bandwidth Environments

This section discusses issues to consider in environments with low-bandwidth connections. You can be selective about which types of connectors you choose, how you operate with intermittent dial-up connections, and where you place your global catalog servers and domain controllers.

Connector Performance

Tests have proven that Internet Mail Service in Exchange Server 5.5 can outperform the X.400 connector by as much as 300 percent, mainly because both RPC and X.400 transports rely on stringent call setups, handshaking, and acknowledgements. However, SMTP doesn't carry nearly as much protocol overhead as these other protocols. With the additional performance improvements made in the Exchange 2000 SMTP stack, such as **CHUNKING** and **PIPELINING** commands, the performance gap between SMTP and X.400 has increased even further.

The only real scenario where SMTP can have a negative impact on the throughput of messages is where there are large message attachments to be transferred over very slow (less than 9.6 Kbps) or high-latency links. Because the size of a large Base64-encoded attachment is more than the same attachment transferred as a body part over X.400, you may decide to either deploy X.400 connectors for these types of links or to implement dual connectors. The latter approach gives you the best of both worlds, because small text messages are transferred over fast SMTP, whereas large

messages go over X.400. It's safe to implement parallel connectors between routing groups, because they use link state information to calculate the best route. If both of these connectors rely on the same physical network link, and the network link goes down, messages do not bounce back and forth between the connectors, as they do in Exchange Server 5.5.

Non-Connected Networks

When locations are not directly connected by high-speed networks, you must be sure to configure an intermittent connection. This is especially true when asynchronous dial-up through a modem is the only form of connectivity between locations. Unlike earlier versions of Exchange, Exchange 2000 does not include a Dynamic RAS Connector. Instead, a more efficient routing method uses one of the supplied connectors (Routing Group connector, SMTP connector, or X.400 connector) over an on-demand connection that is supplied by the operating system. This takes advantage of the Routing and Remote Access component supplied with Windows 2000 Server.

Domain Controller and Global Catalog Server Placement

Windows 2000 sites help to define the physical structure of a network. By definition, a Windows 2000 site is a collection of computers with full-time, reliable connections, based on Internet Protocol (IP) subnets. When a change occurs in Active Directory, those sites can be used to control how and when a change is replicated to domain controllers in other sites.

Exchange 2000 servers should be installed as member servers of the domain and not as domain controllers or global catalog servers, mainly because of performance implications.

It is recommended that you deploy at least one global catalog server in each Windows 2000 site. Because Exchange makes heavy use of global catalog servers, it is important to ensure that all Exchange 2000 servers have fast, local access to a global catalog server in the vicinity.

Each time an address book lookup occurs and each time a message is routed, a lookup on the closest global catalog server occurs. As a guideline, for every four Exchange 2000 servers installed, you need at least one global catalog server.

Public Folders

Exchange 2000 supports multiple public folder hierarchies (also called top-level hierarchies and public folder trees), which provide greater administrative control and new uses for public folders. For example, you might create a separate public folder hierarchy to collaborate with external users and to keep that content separate from the default public folder hierarchy. In another example, an additional public folder hierarchy might be created at a remote location for the users there to access data that is relevant only for them.

Each public folder tree stores its data in a single public folder store. Folders in the tree can be replicated to every server in the organization that has a public folder store associated with that public folder tree.

The multiple public folders feature could affect your public folder strategy. While the default public folder hierarchy is created on every public folder server and its list of folders (hierarchy) is always replicated, additional public folder hierarchies affect only the servers on which they are configured. This means a set of local public folders can be created on a subset of servers. These additional public folders do not have to be replicated to every public folder server. You can have additional public folder hierarchies to minimize the overall size of the default public folder hierarchy, to simplify navigation, and to reduce the cost of replicating the default hierarchy.

Public Folder Hierarchies

The default public folder store contains the All Public Folders hierarchy. In System Manager this is listed as Public Folders. This public folder store must be associated with every mailbox store on a server to ensure that the All Public Folders hierarchy is displayed in IMAP and MAPI mail clients. Any additional hierarchies that you create are considered general purpose hierarchies and are accessible from standard Windows applications in which the folders in these hierarchies can be mapped as network drives using the installable file system (IFS), accessed using WebDAV, and accessed by Network News Transfer Protocol (NNTP) clients. You can create these additional public folder hierarchies or trees and use them as file repositories for departments, groups, or projects.

The following table describes which hierarchies are visible to which clients.

Table 17.2 Clients and available public folder hierarchies

Client	Hierarchy
POP	None
IMAP	All Public Folders
MAPI	All Public Folders
NNTP	Internet newsgroups
WebDAV browsers	All Public Folders and other mailboxes and folders
Windows applications	Other mailboxes and folders shared with IFS

MAPI client permissions for the top-level hierarchy are the traditional MAPI permission (for example, Editor and Owner). Client permissions for the general purpose hierarchies are based on Windows access control standards.

In mixed mode you can have only one MAPI public folder hierarchy per organization; however, you can have multiple general purpose hierarchies per organization.

In the MAPI public folder hierarchy, folders are mail-enabled by default, whereas in general purpose hierarchies, folders are not mail-enabled by default.

Configuring Public Folders

You can configure some of the settings that affect public folders by modifying the properties for a public folder or public folder store. Remember that you can modify these items only by using System Manager. The following table describes the configuration options available in the **Public Folders Properties** dialog box.

Table 17.3 Public folder configuration settings

Tab	Available Settings
General	Specify the description of the folder. Also enable read/unread information to speed up client access.
Replication	Specify which servers within the organization contain public folder replicas. Indicate the times at which this public folder is replicated to designated servers throughout the organization.
Limits	Specify age and storage limits and retention length for deleted items.
Details	Type an administrative description of the public folder.
Permissions	Configure the folder, message, directory, and administrator rights on the folder.

Public Folders in Active Directory

You can set up every public folder in a public store to appear as a mail recipient in Active Directory. To accomplish this. you need to mail-enable a public folder.

To mail-enable a public folder

- Right-click the public folder, and then click **Mail-enable**

After you mail-enable a public folder:

- The System Attendant connects to Active Directory and creates an object for the public folder in a container, such as a users container. This container is specified on the **General** tab of properties configuration of the public folder tree and applies to all public folders in the tree.
- A directory entry exists with a name of Folder *Name* + Global Unique Identifier. Users with access to Active Directory can use the mail address properties of the object to send e-mail to the public folder.
- Additional tabs are available for the mail-enabled public folder in Exchange System Manager (or the Exchange Folders snap-in) and in Active Directory Users and Computers. They are E-**mail Addresses, Exchange General,** and **Exchange Advanced.**
- You can configure the public folder to appear in the global address list for clients such as Outlook.

Public Folder Replication

By default, when you create a public folder store, only one copy of the public folder store exists in the organization. A public folder store can exist in an organization either as a single copy or as multiple copies. Multiple copies are known as replicas. Using public folder store replicas provides multiple, redundant information points: creating replicas also provides for load balancing as clients access the data.

You can specify which folderc (or set of folders) in a public folder store are replicated to a particular public folder store. Criteria could be based on size. document type, or last modified time. This enables you to design an intelligent replication model that takes bandwidth availability into account. A single folder (or its subfolders) can have multiple rules.

For example, a public folder store could have three rules:

- Replicate all items below *500* KB every *30* minutes.
- Replicate all Word documents every 4 hours.
- Replicate all documents modified in the last 24 hours at midnight. This rule would replicate any non-Word documents of more than *500* KB, such as a Microsoft PowerPoint presentation.

Public Folder Affinity

When a client attempts to access public folder data, the client must be able to connect to a server that contains a replica of the data. The client attempts to connect to any replica to present the requested data to the user.

To maximize efficiency, the client attempts a connection to servers in the following order:

1. The default public store for the client. The default public store is determined by the configuration of the mailbox store that contains the user's mailbox. If the default public store is not available, the client receives a list of servers that contain a replica.
2. A server to which the client has an existing connection.
3. Each server in the routing group where the client's public folder server resides.
4. If the client cannot connect to any of the servers in the clients' own routing group, the Microsoft Web Storage System instructs the client to attempt connection to other routing groups. The attempts are made in the order of cost; the lowest cost to reach the routing group in which there is a replica is tried first. Each connector has a **Do not allow public folder referrals** check box. If this is selected, that connector has an infinite referral cost and it does not handle referrals.
5. In the case where connections to two or more routing groups have the same cost, the servers containing replicas are pooled together and a server is selected at random.

Address Lists

In Exchange 2000, address lists have replaced Address Book Views in Exchange 5.5 and earlier. In System Manager, you create a container address list and associate a rule with it. Those rules use the LDAP Search Filter syntax as defined in RFC 2254 and are extremely flexible. As an example, to create an address list of all permanent staff (not contractors) in the Marketing department in Toronto, you can create a single container called Marketing with a rule of:

```
(&(mail=*)(department=Marketing)(l=Toronto)(!(Extension-Attribute-3=Contractor)))
```

In addition, you can create filter rules in System Manager. System Manager allows searching with the rule at creation time to ensure the rule meets your requirements. To ensure that the address list is up-to-date, the Exchange System Attendant polls the directory on a scheduled basis and populates the address lists as necessary.

The default address lists in Exchange 2000 show only mail-enabled objects, but you can change this to show objects that are not mail-enabled. such as resources. The default views are described in Table 17.3.

Table 17.4 Default address book views

Address List Name	RFC 2245 Syntax
Global address list	(&(|(mail=*)(proxyAddresses=*)(textEncodedORAddress=*)) (|(objectCategory=person)(objectCategory=group) (objectCategory=publicFolder)))
All Users	(&(|(mail=*)(proxyAddresses=*)(textEncodedORAddress=*)) (|(&(objectCategory=person)(objectClass=user))))
All Groups	(&(|(mail=*)(proxyAddresses=*)(textEncodedORAddress=*)) (|(objectCategory=group)))
All Contacts	(&(|(mail=*)(proxyAddresses=*)(textEncodedORAddress=*)) (|(&(objectCategory=person)(objectClass=contact))))
Public Folders	(&(|(mail=*)(proxyAddresses=*)(textEncodedORAddress=*)) (|(objectCategory=publicFolder)))
All Conferencing Resources	(msExchResourceGUID=*)

The Address List Service populates the defined address lists by entering the name and location of the list in the showInAddressBook attribute on the user objects in the directory. To view this data, use a utility such as ADSI Edit, which is available on the *Windows* 2000 ***Sewer Resource*** **Kit** companion ***CD.***

Address List Compatibility with Earlier Versions

Be aware that address list updates and new address lists are not replicated back to the earlier Exchange directory service.

Recipient Update Service

Each Exchange 2000 server has the ability to update the address lists by making calls to Wldap32.dll. Only one of these services is active in each Active Direciory domain; the others remain idle. The service is fully integrated with the Exchange System Attendant and contacts a local domain controller to update the address lists based on the rules set in System Manager.

Ideally, the Recipient Update Service runs on an Exchange 2000 server installed in the domain for which the address lists are being updated. If installing an Exchange 2000 server in each domain is not possible, you must create a new Recipient Update Service for each domain that does not have an Exchange 2000 server. For each Recipient Update Service created, an Exchange 2000 server from another domain is selected to run the Recipient Update Service.

CHAPTER 18

External Connectivity

Daniel Martin, Senior Consultant, Microsoft
Jens Trier Rasmussen, Principal Consultant, Microsoft

In today's business climate. it is common to have many types of e-mail systems that must work together. Mergers and group restructuring require support for multiple e-mail systems. Large organizations often have more than five mail systems deployed. Therefore, during a migration period, or due to required applications being on different platforms, it may be necessary to plan for coexistence with varying e-mail systems. Microsoft Exchange 2000 Server provides many connectors to make this task much easier.

When connectors are in place, user experience is seamless across messaging, even if the two systems function differently. The transfer of e-mail and other information between Exchange and the other messaging system is transparent to the user.

This chapter reviews the connectivity tools in Exchange 2000 and provides information about planning to connect to other systems. This chapter explains how to implement these connectors in a production environment.

This chapter also examines each connector and discusses the differences between Exchange 2000 and Exchange 5.5 when using these connectors and connecting to earlier messaging systems such as SNA Distribution System (SNADS) and Professional Office System (PROFS) in an Exchange 2000 deployment.

In This Chapter

Planning and Best Practices

Active Directory and Connectivity

Connector Review

Connectivity in Mixed Mode

Connectivity to PROFS and SNADS

Planning and Best Practices

Before connecting your Exchange 2000 environment to another system, you need to prepare a strategic plan. The following section discusses planning steps to follow for connectors to other e-mail systems and provides a detailed look at the Active Directory directory service and the critical role it plays in Exchange 2000.

Before connecting to another e-mail system, consider the following:

1. Identify the mail systems in this environment.

   Identify each mail system in your organization and the connectors you use to connect to these systems. Exchange 2000 provides direct connectors to the following systems:

   - Microsoft Mail
   - Lotus cc:Mail
   - Lotus Notes
   - Novell GroupWise
   - SNADS (provided in Exchange 5.5)
   - PROFS (provided in Exchange 5.5)

   Exchange 2000 also provides the following connectors:

   - SMTP connector
   - X.400 connector
   - Routing Group connector

   The SMTP Connector can be used to connect to the Internet or other Internet mail systems that support SMTP, (Netscape and OpenMail). The X.400 connector can be used with other organizations that connect through X.400 connectors. The Routing Group connector can connect other Exchange 2000 routing groups within your Exchange organization. All of these join your Exchange organization over the Internet (that is, over a firewall).

2. Determine how many users are on each non-Exchange system.

   This figure will be very important when you design your connector architecture. A large number of users can mean more e-mail traffic. By identifying this early on, you can identify how much mail travels through the connectors and in turn, how many connectors you will need for that system.

3. Document the location of the servers running other e-mail systems.

   After finding the location, determine the type of connectivity used to connect to those systems. Keep the Exchange connector and the system you are connecting to close together. For example, if you have an Exchange server in New York, and a Notes community in Japan, you may decide to install an Exchange server with a Notes connector in Japan.

   Verify that the required network connectivity and protocols are available where you need to connect them.

4. Document the architecture of the other e-mail systems.

   Before you add connectors, verify the architecture behind the other system. For example, a Notes environment may have many domains. Are these domains connected? That is, can you connect to one domain and allow the rest of the mail to route through the Notes infrastructure—or should you add one connector per domain?

   Understand the network that you are connecting to, because the situation may require multiple connectors.

   If your network is similar to the figure below, you have two choices for connectivity. See Figure 18.1.

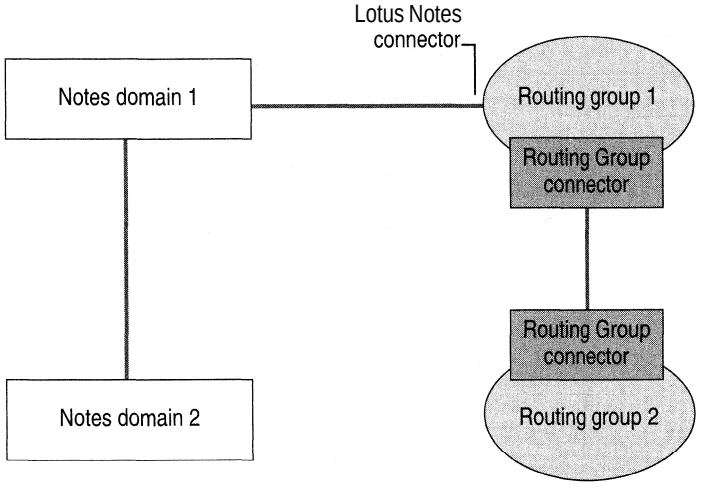


Figure 18.1 Connecting Lotus Notes and Exchange

Figure 18.1 illustrates that with one connector between Notes and Exchange, you can connect the whole Notes organization.

Another choice for connecting other systems is adding a connector from routing group 2 to Notes domain 2. Consider this approach if you are operating in a low-bandwidth environment.

5. Plan your directory synchronization with the other e-mail systems.

   When considering other e-mail system connectivity, look at all the policies and standards of the system to which you will be connecting.

   Know the naming conventions in the Exchange organization and determine whether or not it is the same as that used in the other system. For example, are both systems using a LastName, FirstName convention?

In addition, take into account the implications of deploying Active Directory. Active Directory does not use Directory API (DAPI), but Lightweight Directory Access Protocol (LDAP) as a protocol for directory access. Therefore, if you are using Exchange 5.5, and have built your own directory synchronization scripts, confirm that these continue to work in the Exchange 2000 and Active Directory environments.

If you are synchronizing and connecting with SNADS, you may choose to continue synchronizing directories with Exchange 5.5. Because SNADS and PROFS connectors are not included with Exchange 2000, you need to maintain a portion of your Exchange 5.5 infrastructure to act as connector bridgehead servers to these systems. This also requires directory synchronization between Exchange 5.5 and Active Directory with the Active Directory Connector (ADC).

The Notes and GroupWise connectors have extended directory synchronization (dirsync) functionality with mapping rules for attributes in other e-mail systems. Plan the attributes to include in the dirsync from both systems and, if necessary, define rules for how these attributes appear in their respective systems.

6. Plan Active Directory structure.

   Early planning of the Active Directory structure helps later on with organizing objects from other mail systems in Active Directory. For example, you can store the objects in the same organization unit or keep them in different organizational units. However, how you plan your organizational units is not directly related to your defined mapping rules. Organizational unit layout is only for the ease of administering connected e-mail systems.

   Many decisions in this process depend on how users interact with Active Directory. Specifically, determine if users from other e-mail systems use Active Directory for authentication or just for sending mail. This determines what type of objects to create for mail users of other e-mail systems in Active Directory.

7. Identify applications using the connectors.

   Consider installing a connector to engender co-existence during migration. To do so, determine what type of mail flows through the connectors. Test these applications through the connectors and verify that the e-mail formatting remains the same and that the connectors transfer the necessary types of e-mail.

   Automated applications may create heavy mail traffic. In some situations, you might want to separate the connectors for applications from the connectors used for standard e-mail.

8. Set attachment size limitations.

   Each connector allows limits on the size of documents that the connector can transfer. This may be necessary in low-bandwidth environments to prevent malicious attacks, and when large documents can consume too much CPU processing on the connector server.

9. Determine the connector server hardware requirements.

   In most cases, you must not run connectors on computers that host mailboxes. Before deployment, determine the number of servers and specification for servers you need.

If connectors are underutilized, bundle all the connectors on one server. If one connector is overburdened, you can put this single connector on a dedicated server.

The SMTP connector often runs alone on a single computer. Internet e-mail is often a large proportion of e-mail traffic, so it makes sense to separate the SMTP connectors onto single servers.

For companies of 5.000 users or more, the minimum spccdications for a connector server are a dual 300-MHz processor with 512 megabytes (MB) of RAM and a 1-gigabyte (GB) hard disk.

Although connector server drive space may not be critical, consider increasing the hard disk size if you expect mail to queue on the server. For example, if you connect directly to the Internet without store and forward systems and your link to the Internet is disrupted, mail then queues on the hard disk.

Active Directory and Connectivity

This section describes Active Directory and its connectivity to other mail systems. A recent major change in Exchange 2000 is the tight integration of Active Directory. Active Directory hosts all objects that users can lookup in Active Directory.

This section also illustrates the Recipient Update Service, incorporated into Exchange 2000 to maintain and update the proxy addresses for each user.

Objects

There are differences between an object name in Exchange 5.5 and in Exchange 2000. In Exchange 2000, terminology is shared with Microsoft Windows 2000. Read the list below for objects and to find out how users in other mail systems are stored in Active Directory.

- **User** An account that is created for authentication to a network. It contains a logon password, permissions, and group memberships. A user can have an e-mail address and an Exchange 2000 mailbox.

  For example, when creating a user in Active Directory, you can simultaneously create an Exchange 2000 mailbox for that user. Typically, if Exchange is part of a Windows 2000 infrastructure, the user also has an Exchange mailbox.

- **Contact** A user that can have a target e-mail address on another mail system but cannot log on to the Windows 2000 forest and cannot access resources. A contact in Windows 2000 is equivalent to a custom recipient in earlier versions of Exchange.

For example, if a host (SNADS) user logs on to a host system, and has no need for Windows 2000 forest resources, but does require sending and receiving e-mail to the Exchange organization (within the forest) the user appears as a contact in Active Directory. The contact has a target e-mail address on the SNADS system. Thus, when Exchange users want to send mail to this contact, they can look up the name in the global address list (GAL).

- **Mail-enabled user** A mail-enabled user is a user object with a mailbox on a system that is not Exchange.

  If a Lotus Notes user needs to log on to Windows 2000 to authenticate, the user would appear as a mail-enabled user in Active Directory, but with a Notes e-mail address. Because Windows 2000 authenticates, the user has access to resources within the Windows 2000 forest. Also, because the user is in Active Directory, Exchange users can look up the user in the GAL.

  The difference between a contact and a mail-enabled user is that the contact cannot log on to the Windows 2000 forest, whereas the mail-enabled user does log on to the Windows 2000 forest and can be granted access to forest-wide resources.

- **Organizational units** An organizational unit contains objects such as user accounts. You may decide to separate objects that are synchronized from other systems. To do this, create a new organizational unit where all the users from different systems are stored. If you have Exchange, Notes, and PROFS users in your environment, you can create three organizational units that separate users based on where their mailbox resides. Keep in mind that this is not the only way to structure your own organizational units.

Global Catalog

The global catalog in Windows 2000 generates the global address list (GAL) in Exchange 2000. Users running Outlook 2000 directly query a global catalog server to get addressing information. For other clients, Exchange 2000 forwards requests to the global catalog server. The global catalog holds a replica of specified attributes within Active Directory. You need to understand how Active Directory usage affects global catalog server load and how to efficiently deploy your global catalog servers. For more information, see "Deployment Strategies" in this book, and the Windows 2000 Server and Exchange 2000 Server documentation.

Recipient Policies

In Exchange 5.5, user proxy addresses are created on the **Site Addressing** tab in the Exchange Administrator program. In Exchange 2000, you configure user e-mail proxy addresses in **Recipient Policies** in Exchange System Manager.

In addition, functionality enables Exchange secondary proxies to be automatically generated.

Recipient Update Service

Each domain with objects that have Exchange settings requires a server with the Recipient Update Service, which allows proxy requests defined in the recipient policies to be generated for each user. You will need to determine where you want this service to reside.

To reduce strain on large networks. create a separate domain controller dedicated to this task. If your users do not receive proxies, the Recipient Update Service schedule may be set to **Never run.** Establish how often you run this service.

In the midst of a migration, you may want to run Recipient Update Service more often, perhaps every hour. However, if your environment is stable, and user changes are only happening once a day, you could schedule your Recipient Update Service to run shortly after updates apply to Active Directory.

Connector Review

Exchange 5.5 contains a wide range of connectors to other mail systems. Exchange 2000 provides many of the same connectors:

- **Exchange Lotus Notes connector** Provides connectivity and directory synchronization with Lotus Notes environments. Enables send and receive meeting notifications from Notes.
- **Exchange Microsoft Mail connector** Enables connectivity to Microsoft Mail servers for messaging. Directory synchronization is through the Microsoft Mail dirsync protocols. In addition, meeting requests and free/busy information can be synchronized with the Schedule+ Free/Busy Connector.
- **Exchange Lotus cc:Mail connector** Provides connectivity to cc:Mail networks with built-in directory synchronization between cc:Mail and Active Directory.
- **Exchange Novell GroupWise connector** Provides connectivity and directory synchronization between GroupWise and Active Directory. It supports calendaring functionality, such as meeting requests.

In Exchange 2000, you use Microsoft Management Console (MMC) to configure and maintain the connector queues.

Exchange Lotus Notes Connector

The Exchange Lotus Notes connector can connect Exchange with Notes servers. The connector has built-in messaging, directory synchronization, and meeting request capabilities.

Messaging

The Notes connector retains high fidelity between Notes and Exchange. As in Exchange 5.5, the Notes connector supports features such as:

- Delivery status
  - Read receipts
  - Delivery receipts
  - Delivery Notification Status, NDR
- Options
  - Importance (High, Normal, Low)
  - Type (Private, Confidential)
- Formatting
  - Rich text
  - Doc-links (RTF, URL, OLE)

In addition, the Notes connector allows synchronization between Active Directory and Notes Address Books.

Directory Synchronization

Use MMC to configure the connector for directory synchronization and messaging. Mapping files (tables) with default mapping rules for your attributes are in the following locations:

```
\exchsrvr\conndata\dxamex
\exchsrvr\conndata\dxanotes
```

The Notes connector uses both directories during directory synchronization. The following table lists the directory files and their uses.

Table 18.1 Notes connector files

Directory	File Name	Purpose
dxamex	Mapnotes.tbl	Mapping rules: entries going from Notes to Exchange.
dxamex	Amap.tbl	Defines Exchange fields to be synchronized.
dxanotes	Mapmex.tbl	Mapping rules: entries going from Exchange to Notes.
dxanotes	Amap.tbl	Defines Notes fields to be synchronized.

The files in Table 18.1 are for directory synchronization. These files have default values, and are sufficient in most cases. However, you may need to edit these files to create custom rules.

The Mapnotes.tbl and Mapmex.tbl contain macro language code to create mapping rules. The language contains functions such as ISEQUAL() and TRIM() to operate on the data received from other directories.

Configuration

Install the Notes client version 4.6, or higher. Verify that the Notes version corresponds to your Exchange server. After installation, the connector and the client can use the same DLLs to connect to the Notes server.

Prior to starting the connection, create an ID file for the connector in your Notes environment. This ID allows the connector to connect to the Notes environment. The connector supports passwords and it is recommended that you ure one for security. Also, give this Notes ID particular rights to certain databases.

The following section discusses the configuration of the Lotus Notes connector by examining each tab on the Notes connector property page—accessed using MMC.

General tab

- **Notes Server** Enter the fully qualified domain name (FQDN) of the Notes server to which you will be connecting. An FQDN is a Domain Name System (DNS) domain name that indicates the server's location in the domain namespace tree. It consists of a host name and a domain name, including the top-level domain. For example, host.nwtraders.microsoft.com is a fully qualified domain name. Host is the host name, nwtraders is the third-level domain, microsoft is the second-level domain, and com is the top-level domain. An FQDN always starts with a host name and ends with the top-level domain name.
- **Notes INI file location** Install the Notes client on the Exchange connector server. The installation specifies the Notes client configuration file location. By default, this file exists in *%system%\notes.ini*
- **Connector mailbox** This is a *Mail Router* database associated with the Exchange 2000 domain configured in Notes. This is where the Notes connector delivers all of the mail messages sent to Exchange.

  Note The Exchange connector ID in Notes requires Manager rights with delete permissions on this database.

- **Polling interval** The default is 15 seconds. The connector looks inside the connector mailbox every 15 seconds to confirm that new mail is waiting to bc delivered to Exchange. The 60-second polling interval is not likely to relieve much load on the server, so you can improve user experience by choosing the shorter interval.

- **Notes server language** Select the Notes Server language installed on the Notes Server to which you are connecting.
- **Convert Notes doc-link** Select the type of doc-links transferring to Exchange when a Notes user sends a doc-link to an Exchange user. This setting affects everyone using this connector for Notes and Exchange mail transfers. It is a global setting that cannot be set on an individual basis.
  - **RTF** A doc-link appears to the Exchange user as an attached Rich Text Format (RTF) message.
  - URL If the Notes Server is also a Domino server, it is possible to open the doc-link and its associated documents in a Web browser. The Exchange user receives a URL to the doc-link location.
  - **OLE** If all users have the Notes client on their desktop, you can select this option. When a doc-link goes to Exchange users, the users can double-click the attached icon in Outlook, the Notes client then launches, and the user sees the doc-link on your Notes Server.

Address space tab This is a tab common to all connectors. The address space sets the type of addresses for which the connector is responsible. Most of the organization uses wildcards so that the Notes connector routes any messages that appear as:

NOTES:*@* to **the Notes Connector.**

During your planning, you should decide if you want to deploy more than one Notes connector. With the appropriate address space configuration, you can determine which connector handles which mail. If you set the address space to NOTES: *@ DOMAIN1, only messages addressed to DOMAIN1 will go to this connector.

You can also define connector scope (a feature available in all connectors to other systems). This determines whether the entire organization or just the routing group in which the connector resides uses the connector.

Delivery restrictions tab This tab is common to all connectors and restricts connector usage.

Dirsync tab Use this tab to configure dirsync details. How you configure this tab determines your address book behavior.

- **Schedule** Before connecting to the Notes group, define the dirsync schedule. For most organizations, you can set synchronization to occur only once daily between Active Directory and the Notes address books. Limit the number of synchronization cycles to the minimum necessary. If you update directories once a day, there is no need to run dirsync every hour.
- **Address Book Settings** Address book settings allow you to configure updates for multiple address books. For example, you may enter and manage all the Exchange user—in a different Name and Address Book than other Notes users. Thereafter, your target address book entry differs from your source address book.

By default, both the target and source address book are *names.nsf*: the default Notes Name and Address Rook. You can change these to *source=names.nsf* and *target=foreign.nsf*. This allows the separation of both communities into separate books. If you do this, change the Notes.ini file to add the foreign.nsf to the list of databases that Notes searches when looking for names.

Specify the proper rights to the Connector Notes ID that performs maintenance on the address book. For the target address book, you need manager rights, and for the source address book, read rights. If you use one address book for both the source and target, then grant manager rights on that address book.

Specify which group (if any) needs to be synchronized with Exchange.

- **Synchronization** The **Full Reloads** and **Updates Only** buttons immediately initiate your choice of full reload, or a partial update. They do not change the state of subsequent dirsync cycles. A scheduled dirsync performs a bi-directional partial update.

Import container tab Select where Active Directory stores the Notes entries. Plan your Active Directory strategy before configuring this tab. You can separate each group of Notes users by organizational unit. For example, you can create an organizational unit under Users called *Notes Users* and select that as your container for Notes users.

The Irnport tab specifies how Notes users import. There are three options:

- **Contacts** Users that have e-mail addresses but who do not log on to the Windows 2000 system. These users have no access to Windows 2000 resources.
- **Mail-enabled users** Users that are mail-enabled and who log on to Windows 2000 to access its resources.
- **Disabled Windows account** A disabled account is useful if the Notes user currently logs on to a Microsoft Windows NT® 4.0 domain or will migrate to Windows 2000 later.

Export container tab This extracts the Active Directory objects to be transferred into the target Name and Address Book. Add as many containers as necessary on this tab. Here you can specify contact and distribution list synchronization with Notes.

Advanced tab

- **Notes letterhead** Notes allows you to receive messages with a different template. Choose which letterhead appears on the Notes message when an Exchange user sends mail to Notes.
- **Notes router mailbox** Exchange mail arrives in the Notes router mailbox. The Connector Notes ID requires depositor rights to this database in Notes. By default, the Notes router mailbox is called *mail.box* on most Notes installations.
- **Delivery order** The delivery-order setting default is **Priority:** the message with the highest priority arrives first.
- **Notes database maintenance** The default is **Never-run.** Specify a schedule if you want the connector to handle Notes Database compaction. Select a time when the connector is not busy.

- **Routable domains** If you set up your connector to attach to a hub domain that handles connections to those downstream, indicate all the downstream domains in the Routable domains table, thereby directing mail to travel to these domains.
- **Message size** You can restrict what message size passes through the connector to the Notes users.

Exchange Microsoft Mail Connector

The Exchange Microsoft Mail connector delivers mail between Exchange and Microsoft Mail. Directory synchronization for Microsoft Mail occurs through the Microsoft Mail dirsync protocol of requestor and server post offices. The directory synchronization services and connector are installed simultaneously. You can configure the services through MMC, but remember that there is always a Local DXA Service installed for every instance of the Microsoft Mail connector. The administrator configures the Local DXA service to be a server or a requestor based on topology.

Run Microsoft Mail version 3.2 or Microsoft Mail 3.5 for the Exchange Microsoft Mail connector.

Messaging

The Microsoft Mail connector directory on the Exchange server is in \Exchsrvr\Connect\Msmcon. You can configure the Microsoft Mail connector with MMC.

Configure the connector with a (PC) MTA service. Do not confuse this with Exchange MTA stacks, which Exchange also relies on for mail delivery. The MTA service connects a number of Microsoft Mail post offices. It is therefore important to plan how many MTA services you need and the number of post offices that are under each MTA service.

The Microsoft Mail connector allows mail to go between the two systems, provided users know the fully defined address. The connector installs Exchange Microsoft Mail Post Office on the Exchange server that can interact with other Microsoft Mail post offices.

Directory Synchronization

The directory synchronization portion of Microsoft Mail is done through directory synchronization server and directory synchronization requestor. A directory synchronization server sends updates to all other post offices. There is one directory synchronization server per Microsoft Mail network. A requestor solely sends updates to the server, and receives updates from the server. There is a requestor on each post office in the Microsoft Mail network, including the directory synchronization server.

Note The directory synchronization server is also a requestor.

To make Exchange 2000 a directory synchronization server or requestor, open System Manager, right-click **Connectors,** point to **New** and then click **Dirsync Requestor** or **Dirsync Server.**

When the Microsoft Mail connector is selected frotn Setup, the Microsoft Mail connector and the Free/Busy Schedule connector install by default. Since the dirsync service appears at setup, configure only the connectors.

You must configure the connectors in the following order:

1. **Microsoft Mail connector** Sends mail between Exchange and Microsoft Mail.
2. **Dirsync Requestor or Server** Populates the user list with the other system's mail addresses
3. **Free/Busy connector** Free and Busy schedules appear in meeting requests.

The dirsync service populates the mail addresses and allows for easier selection of the relevant e-mail addresses from either system. The dirsync service can be a requestor or a server. Your current Microsoft Mail configuration determines the required type of service.

Generally, if you have more than one Microsoft Mail post office, you have a Microsoft Mail native dirsync server that handles all of the global address lists. Thus, the Exchange server becomes a dirsync requestor in the native Microsoft Mail Setup, as shown in Figure 18.2.

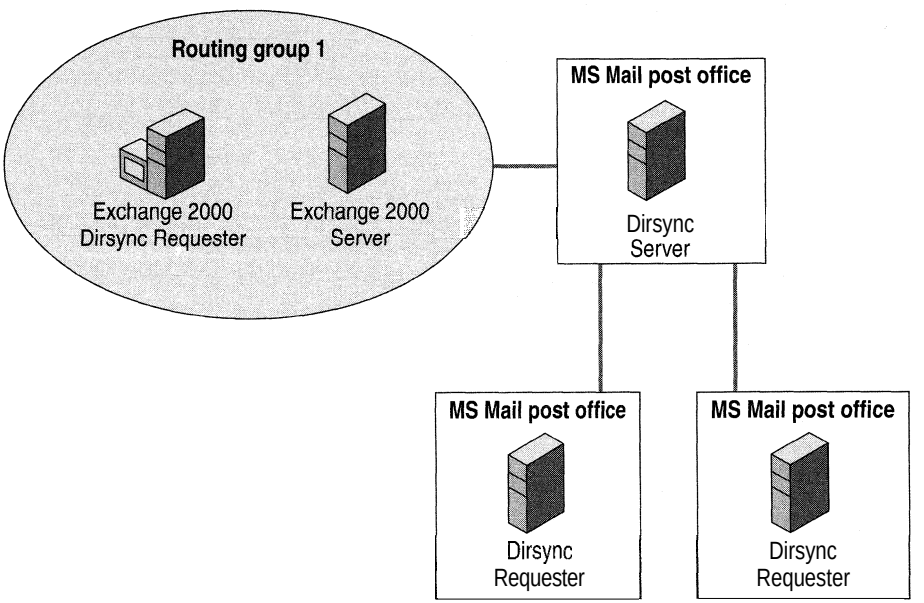


Figure 18.2 Microsoft Mail and Exchange coexistence

In an environment where one Microsoft Mail server runs, it Is possible for it to co-exist with the Exchange server acting as a dirsync server.

Plan which post offices are to be dirsync requestors and which server is to be the dirsync server. In a new environment, make the dirsync server the Exchange server. Remember, there is only one dirsync server in a Microsoft Mail environment. This rule also applies to a company using Exchange 2000.

Microsoft Mail uses a concept of time points T1, T2, and T3 to define its directory synchronization cycle.

- **T1** All the requestors send their updates to the dirsync server.
- **T2** All the updates process on the server. T2 indicates when all changes propagate back to the requestors.
- **T3** Requestors receive the new updates from the dirsync server.

If the directory synchronization is not working after the T1, T2, and T3 time points, you can manually synchronize on the dirsync server and use the command prompt to force the updates. If the dirsync service fails, check the configuration. Manual updates do not necessarily fix the failure.

There are DOS commands in the Microsoft Mail Administrator's Handbook. Run these commands on the mail post office dirsync server, assigned to the M drive.

For more information about the dirsync process when Exchange is the requestor server, see the Microsoft Product Support Services Knowledge Base article Q148309, "XFOR: Manual Dir-Sync with Exchange as Dir-Sync Requestor." For more information about the dirsync process when Exchange is the dirsync server. see Knowledge Base article Q147464, "XFOR: Event ID 178 During Directory Synchronization." Dirsync is scheduled to occur once a day. Use these two articles to write short batch files to force the dirsync to occur as often as you want. This saves you time during the testing phase.

Calendaring

The Free/Busy connector allows free and busy periods to transfer between the two systems.

If the Free/Busy connector needs to be set up, then the additional program Adminsch.exe. This must reside on the dirsync server post office. For a description of how to set up the free and busy replication schedule, see Knowledge Base article Q141755, "PC WSPlus: How to Set Up Schedule Distribution."

Configuration

Read this section to learn about the Microsoft Mail connector configuration. Described herein are tabs for Microsoft Mail connector properties, which can be accessed through System Manager.

Interchange tab

- **Administrator's mailbox** Enter the administrator's mailbox to receive system administrator messages. Monitor this mailbox.
- **Primary language for clients** Select the appropriate language.
- **Maximize MS Mail 3.X compatibility** If your Microsoft Mail network consists of 3.x post offices, select this option.

Address space This is a common tab for all connectors. This identifies address type for which the connector is responsible. By default, wildcards exist in the address space so that the Microsoft Mail connector routes all messages. The default address space appears as:

`MSMAIL:*\*\* to the MSMail Connector`

Local post office tab When the Microsoft Mail connector is installed, a Microsoft Mail 3.5 shadow post office resides on Exchange. This post office is the Exchange entry point. The connector routes all the mail destined to this address to the Microsoft Mail connector.

As long as Exchange is defined as an external post office, do not edit the routing table. Otherwise, you will need to enter the address space in the routing tables of Microsoft Mail. This is so that a Microsoft Mail user sending to an Exchange user has a route to the Exchange server holding the Microsoft Mail connector.

Connections tab Before connecting to Microsoft Mail, list all Microsoft Mail post offices in the list of connections. You may only need to enter one post office, thereby relying on Microsoft Mail to route to the rest of the post offices in the network. If you wish to rely on Microsoft Mail internal routing, specify each post office on an Exchange 2000 post office. Otherwise, the Microsoft Mail (PC) message transfer agent for Exchange 2000 does not pick up mail for these post offices.

When configuring connections, enter the path of the Microsoft Mail post office. If the post office is visible on the network, this field populates.

This tab can also maintain your Microsoft Mail connectors. Once you establish your connections, click **Queues** to see how much mail is queued to go to Microsoft Mail.

Connector MTAs Use this tab to create an MTA service for the post office that this connector services. It is possible to create more than one service. This tab is also used to specify which post offices are in service and to schedule when traffic travels through this connector. The only scheduling available for external service is how often to pick up mail (for example, every 5 minutes). If you want more control over the schedule, run a dispatch program to start and stop the external service on individual post offices.

Advanced tab

- **Message Size** You can restrict the size of messages that pass through the connector.

Exchange Lotus cc:Mail Connector

The Exchange Connector for Lotus cc:Mail supports DB8 type Lotus cc:Mail post offices. Lotus cc:Mail users join the Active Directory as contacts. mail-enabled users, or disabled users. Directory synchronization is built in to the connector, and only requires the latest versions of Import and Export, which are available from Lotus cc:Mail. All connector configuration is available in System Manager.

The connection is easier and more reliable if you are connecting to a Lotus cc:Mail post office running database version DB8. Verify that the routing Lotus cc:Mail post offices have the latest versions of Import.exe and Export.exe installed. Import.exe and Export.exe are available on the Lotus cc:Mail software disks.

Table 18.2 Lotus cc:Mail name formats

Name	Format
Mailboxes for Lotus cc:Mail users*	*Last name, First name*
Mailboxes for other items, such as, automated batch files.	Must begin with ~
Bulletin boards and Lotus cc:Mail mailing lists	Must begin with #

*Hyphenate compound surnames (Smith-Jones)

Automatic Directory Exchange (ADE) is the synchronization mechanism in Lotus cc:Mail. Successful ADE updates inust occur on a daily basis throughout the Lotus cc:Mail messaging system. All post offices must contain an accurate and up to date Lotus cc:Mail directory. When you install a post office in Lotus cc:Mail, decide if ADE resides on the post office or not. ADE allows directory synchronization between one Lotus cc:Mail post office and another.

The current design of the Lotus cc:Mail installation determines the co-existence phase and the position of the connectors. Most large installations follow a hub and spoke design with central hubs and post offices connected to the hubs. It is possible for the mail post offices to have two or three post offices connected in turn, creating longer spokes. The logical mail flow between these post offices may not be as expected, because point-to-point connections can be created between any of the connected post offices.

It is also common for the Lotus cc:Mail administrator to have amended the directory synchronization process and to have created their own mail synchronization procedure that takes into account world-wide time zones and updates at less busy times.

There can only be one connector that has both message and directory synchronization enabled. In Figure 18.3, the connector in London performs both message and directory synchronization. Depending on the configuration, it may be possible to install connectors that synchronize only messages at major hub sites, as it appears in the Hong Kong SAR location. The existing ADE implementation keeps the Lotus cc:Mail directory up to date using the central ADE server (this happens between Lotus cc:Mail post offices only, and does not appear in Figure 18.3).

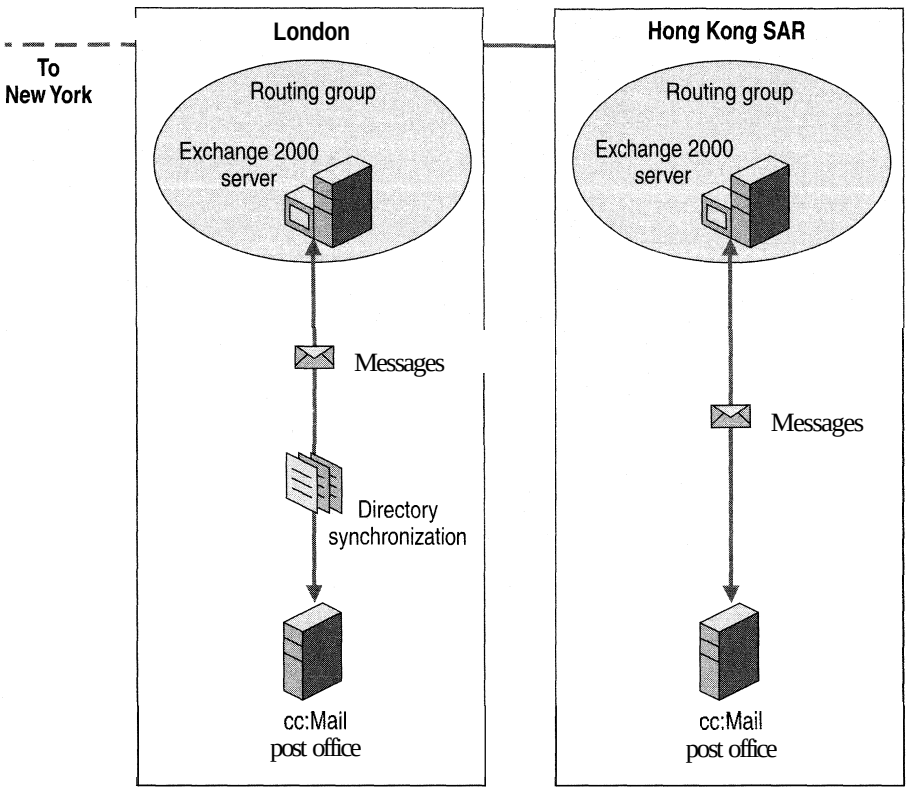


Figure 18.3 Directory synchronization with cc:Mail

Scenario

A typical company will have the Lotus cc:Mail hub locations divided into regions, with Tier 1 hub post offices forming the messaging backbone for the current Lotus cc:Mail messaging system. For example, a company has post offices located in London, New York and Hong Kong SAR. All the company's user post offices connect to one of the hubs for message flow, either directly, or through a Tier 2 routing post office.

The Exchange architecture usually has an identical messaging backbone, with Exchange hub sites in the same location.

There is one Exchange Connector for Lotus cc:Mail at each hub. This will ensure that messages addressed from Exchange recipients to Lotus cc:Mail users (or vice versa) within the same region are not routed across the messaging backbone unnecessarily.

Directory synchronization will be configured at a single hub only.

Recipient containers for all Exchange sites will export from Exchange to Lotus cc:Mail through the *isolation post office*, located on the Exchange Server. The Lotus cc:Mail Router program routes messages between this post office and the Lotus cc:Mail environment, thereby effectively separating Exchange from the Lotus cc:Mail environment. The entire Lotus cc:Mail user directory will import from the isolation post office into a dedicated recipients container in Exchange. An isolation post office is not necessary, but helps to support the connectors. The isolation post office connects to the Exchange connector to exchange mail with Lotus cc:Mail. It is preferable to have no Lotus cc:Mail mailboxes on the Isolation post office and use it solely as a hub to route to other post offices.

Set directory synchronization to occur before and after ADE finishes. The first synchronization ensures that all Exchange directory updates come from the London isolation post office to Lotus cc:Mail. The second synchronization ensures that the Lotus cc:Mail updates received by the isolation post office go to Exchange.

Configuration

The following section lists the configuration of the Exchange Connector for Lotus cc:Mail by defining each tab on the Lotus cc:Mail connector property page. These properties are available in System Manager.

Post Office tab

- **Administrator's mailbox** Enter an administrator to receive messages from the connector for delivery information, and issues that arise with the connector.
- **Lotus cc:Mail post office** Enter the name of the post office you are connecting to the connector. The post office must be a DB8 post office.
- **Path** Enter the path to the post office. You can use a uniform naming convention (UNC) such as *ccmailserver**ccmdata*.
- **Connect as** Enter a Windows account with access to the shared path mentioned above.
- **Allow ADE to propagate synchronized entries** Select this box for updates sent from Active Directory to Lotus cc:Mail to be propagated to other connected Lotus cc:Mail post offices with ADE.
- **Preserve forwarding history** Select this box to preserve the forwarding history of messages sent from Lotus cc:Mail to Exchange. This history appears in Exchange as an attachment.

Advanced tab

- **Message size** You can restrict the size of messages that go through the connector.
- **Exchange-Lotus cc:Mail directory update schedule** Select the time at which you want to run directory synchronization. This should be coordinated with the times set for the Lotus cc:Mail environment. Before synchronizing the directories, verify that all the changes in the Lotus cc:Mail environment have arrived at the Lotus cc:Mail post office where the connector runs. In addition, you need to make sure that the updates inserted in the Lotus cc:Mail post office from Exchange propagate back to connected Lotus cc:Mail post offices.

Address space tab This tab determines the type of addresses for which the connector is responsible. Two fields are required for address space configuration. mailbox and post office. Most companies use wildcards for this so that the Exchange Connector for Lotus cc:Mail routes any messages, such as:

```
at * to the cc:Mail connector.
```

Some companies may decide to separate their traffic across many connectors. You can decide which connector is responsible for which post office by using the address space. The format is:

```
CCMAIL:* at postoffice
```

Import container tab Here you select where the Lotus cc:Mail entries will be stored in Active Directory. Plan your Active Directory strategy before configuring this tab. You can separate each outside group by organizational unit. For example, you could create an organizational unit under Users for "Lotus cc:Mail Users" and select that as your container for Lotus cc:Mail users.

Additionally, this tab specifies how Lotus cc:Mail users import. There are three options:

- **Contacts** Users with e-mail addresses that do not log on to Windows 2000, and therefore, do not have access to Windows 2000 resources.
- **Mail-enabled users** Users that are mail-enabled and that log on to Windows 2000 for access to Windows 2000 resources.
- **Disabled Windows account** Notes Users that might or might not log on to a Windows NT 4.0 domain, but are scheduled to migrate to Windows 2000 later.

Decide which type of address goes in to the directory synchronization. You have the option of entering filters that allow only certain addresses to be propagated to Exchange:

Import all directory entries Select this if you do not want to filter any entries.

Only import directory entries of these formats Select this if you want to receive only certain post offices. For example, if you only want users on PO1 and PO2 synchronize back to Exchange, enter the following two filters:

```
* at PO1
* at PO2
```

Do not import directory entries of these formats Select this to exclude some post offices. For example, if you do not want to receive any users in the post office London, you would enter the following filter:

```
* at London
```

Export container This extracts Active Directory objects to be transferred into Lotus cc:Mail. You can add as many containers as necessary here. Decide at this point if you want to synchronize contacts and distribution lists back to Lotus cc:Mail.

Delivery restrictions tab This is common in all connectors and restricts which users have access to and receive messages from the connector.

Exchange Connector for Novell GroupWise

The Exchange Connector for Novell GroupWise allows mail and directory synchronization between Novell GroupWise and Exchange.

Messaging

A message sent from Exchange to GroupWise causes a lookup in Active Directory to find the GroupWise user (either a contact or mail-enabled user). A GroupWise user in Active Directory has a target address pointing to the Exchange Connector for GroupWise. The connector transforms the message in a format understood by the GroupWise API Gateway; the connector transforms the message into GroupWise API format. The message then goes to the recipient through the GroupWise environment.

A message sent from GroupWise to Exchange causes GroupWise to check the recipients against its link configuration table. In the case of Exchange recipients, this is the application programming interface (API) gateway. The message goes to the GroupWise API gateway and then to the Exchange Connector for Novell GroupWise. The connector delivers the message to the Exchange MTA, and it travels normally to the intended recipient.

The connector converts meeting requests to calendar items on GroupWise; in addition, other types of messages, such as phone messages, convert to e-mail messages on Exchange.

Directory Synchronization

The configuration file for the directory synchronization portion of the connector appears on the connector server in the following directory: \exhchsrvr\conndata\dxagwise.

The following table lists the directory files in this directory and their uses.

Table 18.3 GroupWise configuration files

Directory	FileName	Purpose
Dxagwise	Mexamap.tbl	Defines Exchange fields to be synchronized.
Dxagwise	GwAmap.tbl	Defines GroupWise fields to be synchronized.
Dxagwise	Mapmex.tbl	Mapping rules: entries going from Exchange to Groupwise.
Dxagwise	Mapgwise.tbl	Mapping rules: entries going from Groupwise to Exchange.

Configuration

The following section contains the configuration of the Exchange Connector for Novell GroupWise by examining the GroupWise properties in System Manager.

General tab

- **API gateway path** Enter the path where the gateway is
- **Netware account** Enter a Netware account with rights to the path specified in the API **gateway path.** Make this account a member of the NDS group called NTGateway.
- **Message size** You can restrict message sizes passing through the connector.
- **Delivery order** This defines the order in which mail messages arrive. The default is to deliver by priority.

Address space tab This is a common tab for all connectors. An address space sets the addresses type for which the connector is responsible for routing.

Delivery Restrictions tab This is also a common tab in all connectors. It restricts which users have access to and receive messages from the connector.

Directory synchronization schedule tab This tab enables you to configure the details of your directory synchronization.

- **Exchange-Groupwise directory update schedule** Before connecting to the GroupWise system, define the directory synchronization schedule. For most companies, changes to the address books or Active Directory occur once a day. Decide on how often synchronizing directories will occur. It is advisable to limit the number of synchronization cycles to the minimum necessary.
- **Synchronization** The **Full Reloads** and **Updates Only** buttons initiate your choice of full reload, or a partial update immediately. They do not change the state of subsequent directory synchronization cycles. A scheduled synchronization performs a bi-directional partial update.

Import container tab Select where the GroupWise entries reside in Active Directory. You must plan your Active Directory strategy before you configure this tab. Separate each outside group by organizational unit. For example, you could create an organizational unit under Users for Novell GroupWise and select that as your container for GroupWise users.

Additionally, this tab is used to specify how GroupWise users will be imported. There are three options:

- **Contacts** Users that have e-mail addresses but do not log on to the Windows 2000 system; they therefore, do not have access to Windows 2000 resources.
- **Mail-enabled users** Users that are mail-enabled and that log on to Windows 2000 to use Windows 2000 resources.
- **Disabled Windows account** Notes user that might or might not log on to a Windows NT 4.0 domain, but is scheduled to be migrated to Windows 2000 later.

Filtering tab You can set options for filtering GroupWise directory entries during import. When you do so, specify domains, post offices, or individual recipients included in, or excluded from the import container. Indicate a filter using the same format as a GroupWise address:

```
GroupWise_domain.Post_office.Object_ID
```

You can use the wildcard characters asterisk (*) and question mark (?), where an asterisk denotes any number of characters of any value, and a question mark denotes one character of any value. Table 18.4 has examples of filters to configure.

Table 18.4 GroupWise filter examples

Import filter	Description
NewYork01.PO1.*	Imports only the recipients at post office, PO1, in the NewYork01 domain.
.PO1.	Imports all recipients in any PO1 post office in any GroupWise domain.
.PO?.	Imports all recipients in any GroupWise post office that has a three-character name, where the first two characters are "PO".

Export container tab This is the container that will be used to extract Active Directory objects for transfer into GroupWise. You can add as many containers as needed. You can also synchronize contacts and distribution lists with GroupWise.

Connectivity in Mixed Mode

Exchange 2000 can coexist with Exchange 5.5. If Exchange 5.5 connectors already exist in the environment, you may decide to leave these connectors in place and start adding Exchange 2000 servers to the environment. This would create a mixed-mode organization, with both Exchange 2000 and Exchange 5.5 present. When adding Exchange 2000, bring the ADC into the environment.

Active Directory Connector

The Active Directory Connector (ADC) allows Exchange 5.5 users and mailboxes to synchronize with Active Directory. Requests and lookups pass on to Active Directory. In addition, mailboxes created on Exchange 2000 servers synchronize and appear in the Exchange 5.5 global address list (GAL). Proxy addresses for each user populate and propagate to each platform. The recipient update service updates the recipient proxy addresses, thereby providing users the appropriate proxies for their Active Directory entries. The ADC updates the following objects into the Active Directory from Exchange 5.5:

- Exchange 5.5 custom recipients become contacts in Active Directory.
- Exchange 5.5 mailboxes become mailbox-enabled users with a Windows 2000 account or a disabled Windows 2000 account.

For more information on the ADC, see "Active Directory Integration and Replication" in this book and the Exchange 2000 Server documentation.

Routing in Mixed Environment

When you introduce Exchange 2000 servers to your environment, identify the current connectors used on Exchange 5.5. If you need PROFS, SNADS. or calendar connectors that are not available in Exchange 2000, you need to keep Exchange 5.5 servers in the environment. For improved message throughput, consider keeping the connectors on a separate server in the same site. If you place connectors in a separate routing group, you defeat the purpose of properly deploying connectors in the presence of WAN links.

You can deploy Exchange 2000 and Exchange 5.5 servers in the same routing group with no connector bottlenecks; this functionality is supported in Exchange 2000.

Connector Migration

As mentioned earlier, configuring the connectors on Exchange 2000 is very similar to configuring connectors on Exchange 5.5. The following procedure outlines a high-level process for introducing the new connectors in Exchange 2000 and removing them from Exchange 5.5.

1. Install the required connectors to other systems on Exchange 2000
2. Create an SMTP connector on the Exchange 2000 server.
3. Create an Internet Mail Service or add an address space to the Exchange 5.5 server.
4. Verify that the ADC is synchronizing between sites.
5. Remove the connector from the Exchange 5.5 site.

Connectivity to PROFS and SNADS

The Exchange PROFS connector and the Exchange SNADS connector are not included in Exchange 2000. Therefore, these connectors must remain on Exchange 5.5 servers.

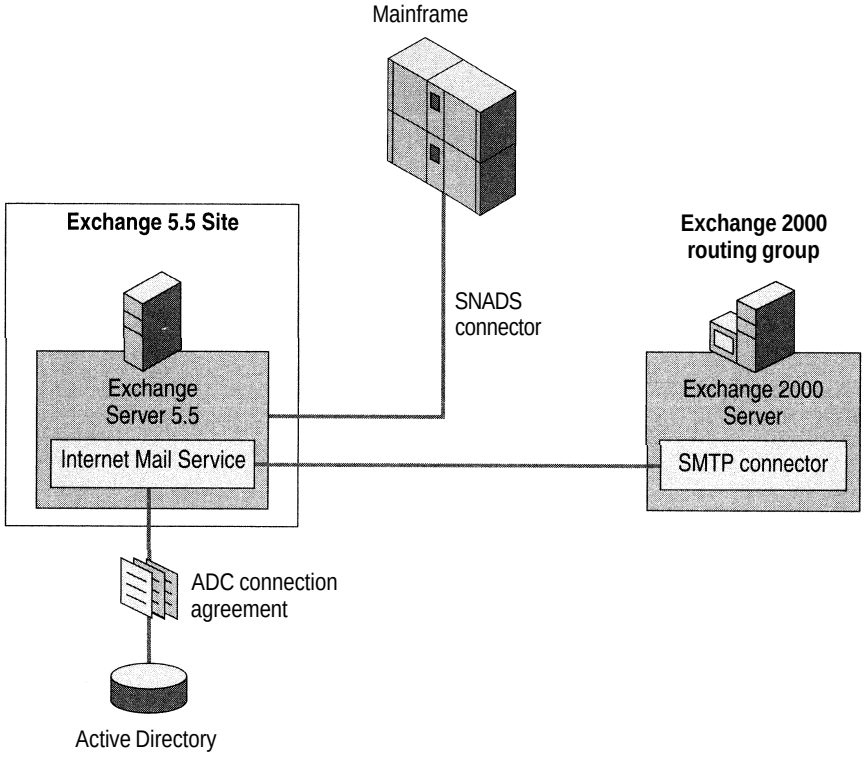


Figure 18.4 Exchange 5.5 as a SNADS connector gateway

Figure 18.4 illustrates the basic configuration for connecting to PROFS and SNADS systems in a mixed-Exchange environment. This allows the generation of PROFS or SNADS addresses for contacts and mailboxes. Since the templates and the proxy address generators are installed on Exchange 5.5, you need to replicate these templates to Active Directory. The ADC will allow for synchronization of those proxies across platforms to the Active Directory.

Directory Synchronization with SNADS and PROFS Connectors

As discussed in the planning phase of this chapter, directory synchronization is critical for getting mail to flow between disparate systemr With SNADS and PROFS, the connectors remain on Exchange 5.5 aerverr. In that case, if you already hate a SNADS or PROFS solution on Exchange 5.5, then continue to ure the synchronization process established for that environment, even as you migrate to Exchange 2000. It is important to know that the entries continue to synchronize on the Exchange 5 5 server. The entries that do this from the host system into Exchange 5.5 and are then synchronized with Active Directory by the ADC.

Microsoft Metadirectory Services

If you have multiple directories from various systems, you may consider using Microsoft Metadirectory Services. With Metadirectory Services, you can receive and send updates from many systems, including flat files created and understood by many host systems (PROFS, for instance). In addition, Metadirectory Service can synchronize directly with Exchange 5.5 or Active Directory. If you decide to incorporate Active Directory, Metadirectory Services can help bring all of your directories into Active Directory. Metadirectory Services can also synchronize with other e-mail systems. For more information about Microsoft Metadirectory Services, see "Inter-Organization Replication and Directory Synchronization" in this book.

PART 5

Advanced Deployment Planning

In This Part

Chat and Instant Messaging Services

Inter-Organization Replication and Directory Synchronization

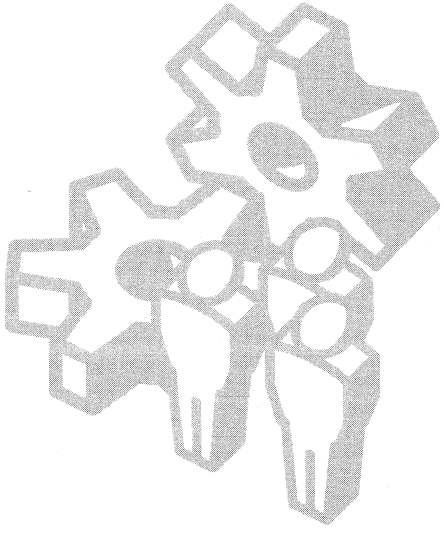
Branch Office Scenarios

Corporate Backbone Scenarios

Hosted Service Environments

Security Sensitive Environments

Outlook Web Access



CHAPTER 19

Chat and Instant Messaging Services

Michael Aday, Senior Consultant, Microsoft

Microsoft Exchange 2000 Server introduces real-time collaboration services that allow distant users to collaborate, to be better connected, and be more accessible for business communication. New services include data conferencing and video conferencing (which are shipped as a separate product), and Instant Messaging. These services, along with Chat Service (introduced with Microsoft Exchange Server version 5.5), can be deployed in various ways. This chapter presents information about best practices and initial approaches to deploying Chat Service and Instant Messaging in Exchange native-mode and mixed-mode environments.

The recommendations provided here are guidelines for including the Exchange 2000 real-time collaboration services in a messaging infrastructure. These recommendations might not be applicable in all customer scenarios. Examine your own implementation plans before following these recommendations.

In This Chapter

Chat Service Overview

Deploying Chat Service

Instant Messaging Overview

Deploying Instant Messaging

Chat Service Overview

Internet Relay Chat (IRC) has existed for many years as a collaborative capability hosted by larger Internet Service Providers (ISPs). With the release of Microsoft Exchange 5.5, you could add this collaborative service to an existing Exchange messaging infrastructure to provide a virtual meeting place for scheduled or spontaneous text-based discussions. Using Chat Service allows a group of people to share ideas in a textual format, in either moderated or unmoderated real-time conversations.

In Exchange 2000, Chat Service takes advantage of Active Directory and its inherent management benefits, while still adhering to IRC standards. Chat Service uses Active Directory to store configuration information and objects such as chat communities, channels, classes, bans, and chat servers. The Exchange 2000 System Manager snap-in can consolidate these items into one interface. In addition to storing configuration and object information in Active Directory, Chat Service uses authentication and access controls from Microsoft Windows 2000 to determine a client's ability to connect to a chat community or join a chat channel, and to define administrative roles such as system operators and chat administrators.

The Exchange 2000 Chat Service implements the IRC (RFC 1459) and Extended IRC (IRCX, currently in RFC draft) protocol standards, allowing for the use of IRC clients such as Microsoft Chat Service. Chat Service allows you to create channels for one-to-many and many-to-many text conversation. Chat administrators can moderate the use of and access to chat communities with bans and classes, and allow users to host or inoderate a chat channel's content.

Chat services in Exchange 2000 have been designed to support greater scalability than earlier releases. While Exchange 5.5 Chat Service relies on portals between chat servers to add additional concurrent connections to the chat infrastructure, Exchange 2000 can host an equivalent number of users on a single server. It is possible to consolidate larger chat channels hosted on several Exchange 5.5 servers onto one Exchange 2000 server.

Deploying Chat Service

It is easier to manage Chat Service when the information is stored in Active Directory. To do so. Active Directory must first be deployed in the organization. In addition, a Windows 2000–based server cluster must be available to host the service to utilize Chat Service's clustering capabilities. While clustering will not preserve specific chat server channel content, it will provide for failover of persistent channels that have been defined in the chat cluster.

Deployment of Chat Service in Exchange 2000 is fairly straightforward; however, before you deploy Chat Service, you should make some decisions about chat channels. For example, do you need registered (permanent) channels in your environment? What types, and how many channels should you host? Do you want to host channels that are accessible from outside your firewall?

Chat channels can support a variety of channel modes that determine the types of connections and content that can be hosted by the channel. *Auditorium mode*, for example, is used to host a very large number of concurrent users who are listening predominantly to a few presenters. This channel mode prevents each user in the channel from seeing the other users; each user sees only the presenters, while the presenters can see all the members of the channel. For more information about the various channel modes, see the Exchange 2000 Server online documentation.

Multiple channel modes can be set when you create the channel. Administrators and system operators can also change channel modes by using a chat client.

Communities and channels in Chat Service can be protected with a variety of authentication mechanisms. Chat clients can authenticate with Exchange 2000 Chat Service by using the Microsoft Windows 2000 security service provider interface authentication mechanism.

Chat Service Deployment Process

An overview of the Chat Service deployment process follows:

1. Plan the infrastructure, determining server locations and properties of required objects, such as the number and types of channels. Determine channel authentication, define bans to be created, and determine which users will administer the services.
2. Establish a meaningful naming convention for your chat servers. Be aware that a naming standard may already be in use. Where applicable, adhere to the existing standard as closely as possible to prevent confusion.
3. Install servers to host the services. Exchange 2000 Chat Service can be installed as a service on an existing Exchange 2000 server, for example, a server already hosting Instant Messaging or Exchange 2000 Conferencing Services management components. If you want to deploy multiple services on one server, recognize that a heavily-utilized Chat Service requires significant server resources in terms of memory, CPU, and network throughput to meet customer response-time expectations.
4. Configure objects within the chat topology. If you do not want any registered channels, bans, or classes, no objects need to be created.
5. Create the required Domain Name System (DNS) resource records (typically, a host record) for each chat server. Decide which servers, if any, will be able to host externally-accessible chat channels. Remember that there may be internal and external DNS requirements; be sure to register the servers in the appropriate zones.
6. Make chat client software available to users.

For more information about the deployment process, see the Exchange 2000 Server online documentation.

Chat Server Scalability

Chat Service scalability has been significantly enhanced when compared with the implementations in Exchange 5.5. For example, Chat Service no longer uses portals. These changes permit up to 20,000 concurrent users per server.

During an IRC session, the server can resolve DNS names for the client upon connection and cache those names for the duration of the session. This procedure is useful for access control and attack protection.

Once connected to the server, users will generate traffic when they post messages or when messages are received. In most cases, these transfers will be small and come in periodically, depending on the activity level of the channel. Users continue to generate network traffic to the server for the duration of the connection.

Chat Server Location

Small companies that are highly centralized can have a single server providing chat services throughout the company.

The chat infrastructure topology in larger companies will likely mirror the distribution of data center hubs in the company's WAN topology. To facilitate localization of network traffic associated with logon and operation of chat clients, larger multinational or geographically-distributed companies may want to locate Exchange 2000 chat servers at the departmental or regional data center level in the company. For larger companies that use Chat Service as a supplemental collaboration component for other event, such as company-wide broadcasts, servers can be centralized in a single location.

For example, a company has approximately 15,000 employees geographically dispersed between three regional data centers (Seattle, Houston, and New York). The New York location has 500 to 1,000 users, the Houston location has 2,000 users; 1,500 users are remote and the balance of the users (10,500) is located in Seattle, the company's headquarters. It will be most beneficial to locate the chat servers in the Seattle location. If additional resources are available, it could be advantageous to place a server in Houston as well. However, if multiple servers are used to host content for a single channel, the company might want to develop a custom IRC client that can post to each instance of the channel on each of the servers. The company could also simply host all 15,000 users on one server in the Seattle location.

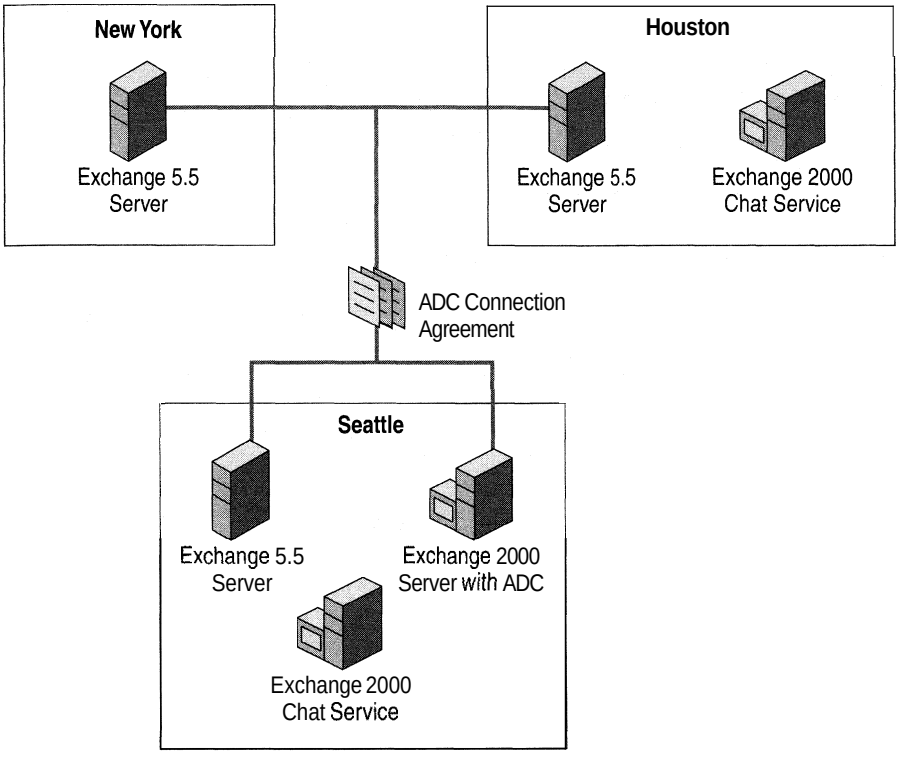


Figure 19.1 Large company deployment with Active Directory Connector (ADC)

If you have previously deployed Exchange 5.5 and want to pilot Exchange 2000 real-time collaboration components such as Instant Messaging and Chat Service, you can do so with minimal impact to the existing deployment by deploying sufficient Active Directory and Exchange 2000 Server resources to host the services. If you support less than 5,000 real-time collaboration users, a single server is sufficient for Exchange 2000 Instant Messaging and Chat Services along with another server that hosts Active Directory and any ADC connection agreements that may be necessary to synchronize user display name and account information between the production and pilot infrastructures. The Active Directory Connector is not required in all configurations, but is necessary in mixed-mode Exchange 5.5 and Exchange 2000 topologies.

As concurrent user connections increase, or when you choose to deploy a more stable pilot, you should host these services on separate computers, beginning with the ADC and then if resources allow, Instant Messaging and Chat Service.

Depending on the concurrent usage of the Chat Service, you might decide to deploy servers that are shared by multiple services. However, you will likely exceed the capacity of a single server sooner or later.

To minimize the administrative overhead associated with the server, you should bundle services that do not save data in Web Storage System, if possible. If you do, the server could host Instant Messaging, Chat Services, video conferencing, and Exchange Collaboration Services management components for as many as 300 users, depending on the amount of activity. This assumes fairly constant data and video conferencing between 2 to 10 concurrent users, and active Instant Messaging users sending three or more messages per hour with 20 or more people on each user's Instant Messaging contact list.

Note that, as with all server-sizing estimates, the above example is a rough estimate based on a hypothetical user load. If you choose to deploy these services on a single server and it fails, all of these services are unavailable. As services grow more crucial, it is important to partition services onto separate servers.

In branch office deployments, the need for Chat Service within a company should be considered independently of WAN topology. If you expect specific client populations within the company to use Chat Service heavily, it might be beneficial to put a local chat server source close to the clients.

Network ddress ranslation and Internal i

IRC services are not affected by network address translation or by firewalls, provided the clients can connect to the ports, as described in the following section.

Security and TCP Ports

Chat Service in Exchange Server 2000 is self-contained. It requires only Active Directory to save the configuration and management information about the chat infrastructure and host servers and to verify access permissions for secured channels. Chat Service transmits information over known ports for IRC (port 6667, by default) and DNS resolution (port 53), simplifying firewall-specific issues.

Clients connecting from remote locations through ISPs can connect to channels on servers offering Chat Service if they are able to resolve external DNS entries. This poses a risk because there is no inherent protection, such as a session layer encryption, of the content of an IRC session even if clients have been authenticated with a strong authentication package such as MD5 Digest or NTLM protocol. During an IRC session, messages are sent back and forth between the client and the server as clear text over port 6667.

Because users can send sensitive information and the information travels as clear text, chat security can be a concern. IRC sessions currently lack encryption support, such as Secure Sockets Layer (SSL) encryption or Transport Layer Security (TLS) encryption. Virtual private network (VPN) protected sessions, such as Point-to-Point Tunneling Protocol (PPTP), Layer 2 Tunneling Protocol (L2TP), or IPSec tunnel (and IPSec arc available in Windows 2000) can be used to encapsulate an entire communications stream; however, if you use network address translation for servers and clients, you could have connectivity problems.

Basic Chat Service Recommendations

The following are guidelines for deploying Chat Service.

- Keep the naming standards for channels and servers as simple as possible. Be aware of the use of internal names, as described earlier in this chapter. Verify that the proper service records have been entered in all internal DNS servers, and if you offer chat services across a public network, in external DNS servers.
- Deploy chat servers to high population client locations based on network topology.

Instant Messaging Overview

Instant Messaging in Exchange 2000 permits small groups of two or more users to exchange basic HTTP messages using a small client-side component and an Internet Server Application Programming Interface (ISAPI) server extension running on a Microsoft Internet Information Services (IIS) server. The main difference between e-mail and instant messages is that instant messages are not saved in Exchange 2000; once a message has disappeared from the screen, it is gone forever. Another difference is that Instant Messaging informs users when their contacts are logged on to the system. This presence information is the most valuable feature of Instant Messaging.

Instant Messaging uses HTTP 1.1 over TCP port 80 and uses a simplified address (based on RFC 822) to identify users. Instant Messaging can be installed separately from other Exchange services; however, it requires the existence of an Exchange 2000 organization.

An Instant Messaging implementation is comprised of the following components:

- **Instant Messaging home servers** These servers are responsible for hosting user accounts and tracking each user's contact lists for posting presence information. Each Instant Messaging home server can support approximately 10,000 online users.
- **Instant Messaging routers** These servers route incoming requests from Instant Messaging clients to their home servers. By placing Instant Messaging Router servers at strategic locations in the deployment architecture, systems can scale more smoothly and provide additional security and management capabilities. Each Instant Messaging router can support approximately 50,000 online users. Routers also allow companies to establish and maintain single Instant Messaging namespaces for all Instant Messaging users, regardless of their home servers.
- **Clients** Computers with Instant Messaging client components installed on them.

- **Windows 2000 domain controllers** Instant Messaging uses NTLM protocol and digest authentication to permit a user access to the service. Windows 2000 domain controllers authenticate users and return configuration information to Instant Messaging Microsoft Management Console (MMC) consoles. Domain controllers also facilitate Instant Messaging servers in determining whether a resource is local or needs to be relayed to some other network location.
- **Firewalls** The presence of firewalls is likely in enterprise-wide Instant Messaging deployments. lnternal firewalls in a departmental deployment are also possible.
- **DNS records** Instant Messaging relies on DNS to return name resolution information for specific services, such as Simple Mail Transfer Protocol (SMTP) and mail exchanger (MX) host records. It is essential to have the correct A (host) and SRV (service) records published to the proper DNS server in an Instant Messaging deployment to ensure that clients are able to resolve host name information properly.

More than one Instant Messaging domain can exist in a deployment. In larger, geographically dispersed companies or in installations that are supported by more than one Windows 2000 forest, multiple Instant Messaging domains should be created. These scenarios are explored in greater depth later in this chapter.

Deploying Instant Messaging

To deploy lnstant Messaging, you will need Windows 2000 Active Directory and because Instant Messaging is installed within the context of an Exchange 2000 organization, you will also need information about any existing Exchange 2000 topology.

If you are upgrading an Exchange system to Exchange 2000, the installation is relatively straightforward. However, you can pilot Exchange 2000 features such as lnstant Messaging before concluding the upgrade of an existing Microsoft Windows NT 4.0 and Exchange 5.5 environment. All that is required to install Instant Messaging is Windows 2000 Active Directory. It is possible to connect the existing Exchange 5.5 environment and the new Exchange 2000 Instant Messaging installation by using ADC.

The following sections discuss the deployment of Instant Messaging in mixed-mode and native-mode Exchange environments and provide direction on possible implementation considerations based on company requirements and typical installation practices.

Instant Messaging Deployment Process

An overview of the Instant Messaging deployment process follows:

1. Plan the Instant Messaging network, including its topology and server resources. Determine how many home servers and Instant Messaging routers are required to support the users. Instant Messaging home servers host user accounts; Instant Messaging routers route messages to home servers.
2. Establish a naming convention for your Instant Messaging servers and user addresses.
3. Install Instant Messaging on servers running Exchange 2000 Server where you plan to create Instant Messaging home servers and Instant Messaging routers.
4. Configure your firewall topology by defining the local IP networks and proxy server that your Instant Messaging client uses.
5. Use IIS to create a virtual directory for each Instant Messaging home server and Instant Messaging router.
6. Create the required DNS resource records for each home server and Instant Messaging router.
7. Set administrative permissions.
8. Set the password policy on the domain controller. This is necessary only if you use digest authentication. If you use NTLM authentication, a policy change is unnecessary.
9. Enable users to access Instant Messaging. Create new Instant Messaging accounts, or enable existing user accounts for Instant Messaging. If users are not on the system, first create Windows 2000 user accounts.
10. Make the Instant Messaging client software available to users in your company.

For more information about the Instant Messaging deployment process, see the Exchange 2000 Server online documentation.

Exchange 2000-Only Deployment

Instant Messaging is installed as a separate component in Exchange 2000 Setup. If you are installing the service in an existing Exchange 2000 environment, the process is simple. However, the following issues and considerations must be resolved before you deploy Instant Messaging:

- Determine the number and location of Instant Messaging servers. This should be based on the location and number of users in your company.
- Investigate physical LAN and WAN topologies and issues. If your company is multinational or has users that are participating in the services over very slow links, you may want to place servers closer to those groups of users. If your company hosts multiple SMTP domains based on country or region, you might want to set up more servers and Instant Messaging domains to facilitate management or to set up a configuration that closely matches your SMTP architecture.

- Choose the administrative group for each Instant Messaging server. After you place an Exchange 2000 server in an administrative group, you cannot move it to another administrative group.
- Determine if access to the Instant Messaging system from other public networks, such as the Internet, is necessary. Secure access from external networks requires additional servers acting as Instant Messaging routers, HTTP proxy servers, and reverse proxy servers.

Depending on the existing network topology and the intended audience for the services, a wide variety of implementations are possible.

Exchange 5.5 Deployment

Instant Messaging can be deployed as an addition to an existing Exchange 5.5 environment, because the client is a separate software installation. It is not necessary to upgrade in totality to Exchange 2000 for Exchange 5.5 to use Instant Messaging. However, certain requirements must be met to provide this functionality.

Exchange 2000 Instant Messaging in an Exchange 5.5 environment requires that Active Directory and the Windows 2000 infrastructure is deployed. Instant Messaging relies on Active Directory for the mapping of physical URLs to logical URLs and for the initial digest authentication that determines a user's access to the system. **A** Windows 2000 infrastructure can be deployed and connected to the Exchange 5.5 environment by using Active Directory Connector (ADC). You do not need to change the existing Exchange 5.5 environment, provided the target server for the ADC connection agreement is running Exchange 5.5 Service Pack (SPI).

Instant Messaging can be deployed only on a computer running Windows 2000 Server. However, for smaller companies, ADC and Instant Messaging services can be installed on the same server. This ADC installation must synchronize existing display names and user's SMTP addresses from the production infrastructure to the host Windows 2000 environment. Once these accounts have been replicated into Active Directory, they can be Instant Messaging-enabled though the Active Directory Users and Computers snap-in, or you can bulk-enable users through an Active Directory Service Interfaces (ADSI) script.

Once properly installed, Windows 2000–based clients that are using Instant Messaging will display the friendly name rather than the e-mail alias. If Instant Messaging is deployed on earlier Windows clients, you must install extensions that allow the operating system to communicate with Active Directory. For more information about Active Directory Client Extensions, search the Microsoft Web site at http://www.microsoft.com or see the Clients directory on your Windows 2000 Server compact disc.

Instant Messaging Routers and Home Servers

Several functions can be assigned to multiple servers in an Instant Messaging architecture. Specifically, Instant Messaging servers can serve two distinct functions: Instant Messaging router and Instant Messaging home server. These functions can be deployed on the same computer. This is a viable configuration for smaller companies that host a relatively small number of users, have one geographic location, and connect to the Internet using an ISP. Figure 19.2 illustrates a typical Instant Messaging deployment in a small company.

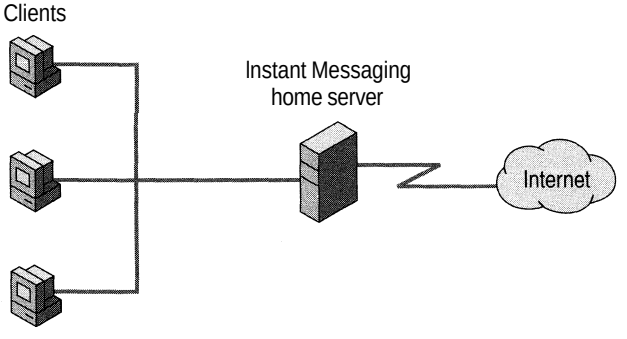


Figure 19.2 Small company deployment

In this type of deployment, one computer performs all of the Instant Messaging functions associated with an Instant Messaging router and home server.

If you deploy more than one Instant Messaging user home server, a separate server must perform the role of the Instant Messaging Router to prevent partitioning of the Instant Messaging namespace and consequent poor performance. Deployment of a separate Instant Messaging router allows incoming connections from Instant Messaging clients to be redirected to their appropriate home server more readily and prevents all Instant Messaging servers from being registered in an externally visible DNS. Figure 19.3 illustrates how a mid-sized company might deploy Instant Messaging.

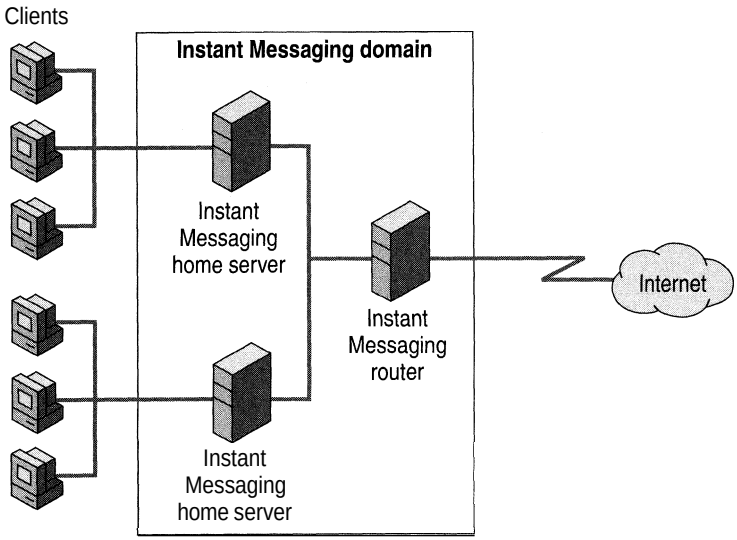


Figure 19.3 Mid-sized company deployment

If your company is likely to exceed the capacity of one server, you will want to have one or more Instant Messaging router servers. The deployment of one or more Instant Messaging router server is necessary in the following instances:

- Your company wants fault tolerance or load balanced Instant Messaging router servers for performance reasons. Instant Messaging routers can load balance by using round-robin DNS or, more preferably, by using network load balancing.

  A company with multiple e-mail domains, such as a multinational company that has partitioned its SMTP namespace geographically. For example, a company has an SMTP namespace, us.contoso.tld, for e-mail within the United States and other SMTP namespaces for other geographic areas such as, gb.contoso.tld. This company would have multiple Instant Messaging domains with one or more Instant Messaging routers associated with each of them.

  > **Note** Throughout this chapter, the root domain name *.tld* is used to distinguish fictitious or sample domain names from actual domain names. Typically, the root domain name is .com, .edu, .org, .net, and so on, rather than .tld.

- A company has a geographically distributed infrastructure. Depending on the available WAN bandwidth, Instant Messaging router servers would either be deployed in areas with Instant Messaging home servers or in a central location, such as the company headquarters. To prevent excessive traffic to the central location, Instant Messaging home servers would be distributed geographically to each location with Instant Messaging users.
- A company has more than one Windows 2000 forest. Each forest is an independent Instant Messaging installation with an associated Instant Messaging domain. This installation is redundant and should be avoided except when the forest namespace is no longer contiguous.
- A company has exceeded 50,000 concurrent users for an Instant Messaging service in a particular location. As an example, a company has 80,000 employees in one geographic location and plans to deploy Instant Messaging services for all 80,000 users.

As organizational topologies grow larger, additional servers might be required to handle external and incoming requests, to host larger numbers of concurrent users, and to route users to their appropriate home servers more effectively. These servers may be deployed in multiple geographic locations or in a central location, depending on a company's needs; however, some of the new server functions are more efficient when placed near a high speed Internet connection.

Companies that want to protect their internal environments more extensively should consider deploying additional servers. The following list contains the components that a company can use to enhance its security:

- **HTTP reverse proxy servers (external)** Used to send all outgoing information from internal Instant Messaging routers to external addresses. You can deploy these servers to protect the internal architecture from public network-connected clients. When deployed, all information appears to be coming from one or small number of servers, based on an HTTP reverse proxy server per Instant Messaging domain. These servers should be placed near high-speed Internet connections or near dial-up services.
- **HTTP proxy sewers (regular outbound)** These servers consolidate outgoing HTTP requests and might already be in place. Instant Messaging uses HTTP requests and requires no special configuration on these servers. However, there is no benefit in caching the returned information because the information is client-specific. These servers should be set up near high-speed Internet connections.
- **Instant Messaging routers for multiple Instant Messaging domains (matched to multiple SMTP domains)** These servers should be distributed to serve larger user populations and consolidated where user population is sparse and in remote locations.
- **Instant Messaging home sewers** These servers can be deployed at remote locations to reduce the long distance bandwidth consumed by user log on and presence update information. This assumes that remote groups work more frequently with one another, therefore presence notification and messaging traffic is confined mostly to those locations.

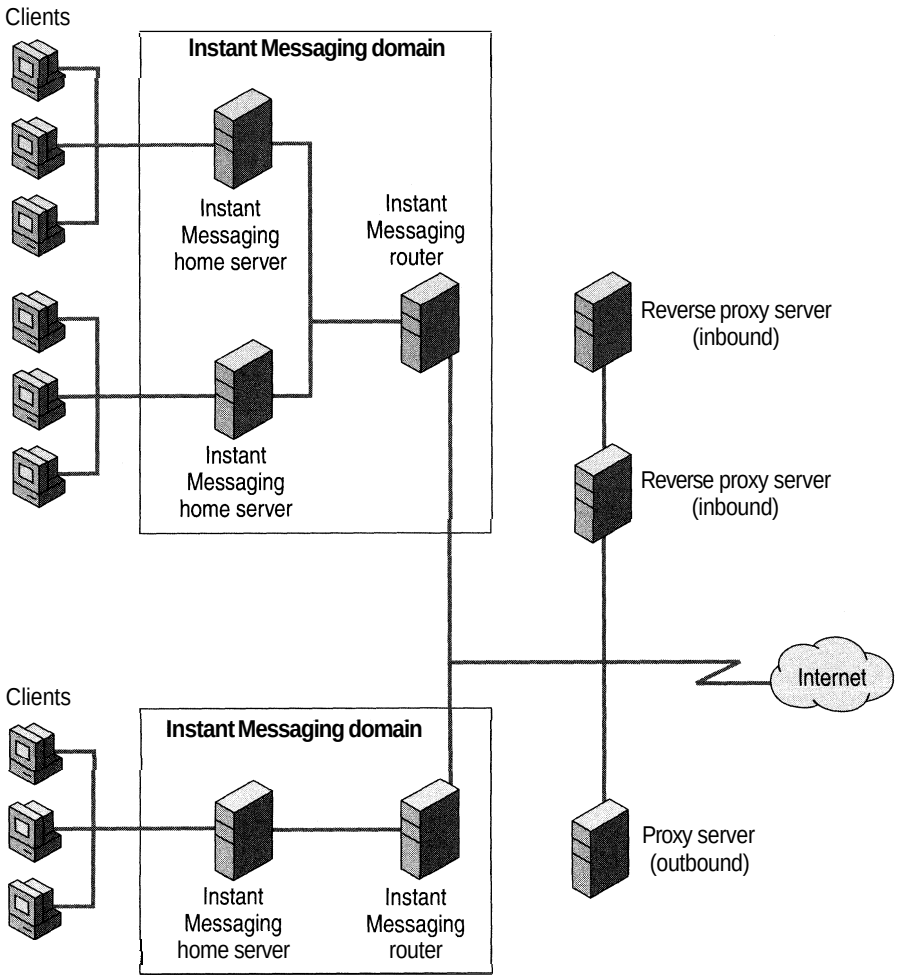


Figure 19.4 Large company secure deployment

Relaying Requests Across the Network

In a global company, there are likely to be many disparate subnets and internal firewalls. Instant Messaging routers determine whether the source server can make the request to the destination directly, whether to forward the request and return the results to the source server, and whether to reject the request. Routing decisions are based on network topology and on source and destination IP addresses of requests, and functions similar to a network basic input/output system (NetBIOS) component on network routers.

By default, information is not published outside the firewall. You must configure the firewall to allow traffic to the Instant Messaging router. You can configure Instant Messaging settings in System Manager in the **Global Settings** dialog box.

i ; Names for Instant :ssaging vers

Instant Messaging clients use HTTP to communicate with the Instant Messaging home server and router server. Microsoft has developed a protocol called RVP that converts the potentially lengthy URL that is formed and used during these communications (http://SRV lookup *result*/instmsg/aliases/*username*) into shorter SMTP-formatted addresses. In many cases, this name is the pre-existing SMTP address of the user or the pre-existing address with an "im" prefix attached to the domain name, such as suzan@im.contoso.tld. Instant Messaging uses service location records in DNS for the RVP services to make this referral. This service registration is similar to MX record registration for SMTP mail hosts in DNS. Records are added so that a client requesting a specific service such as RVP service can identify a host that can provide that service.

Registering Instant Messaging Routers in DNS

Instant Messaging router servers must have the fully qualified domain name (FQDN) of the server (for example, im.contoso.tld) registered in DNS with a SRV (service) resource record.

To create a DNS record for the Instant Messaging routing sewer

1. Open the DNS Manager console and double-click **Forward Lookup Zone.**
2. Right-click the domain to which you want to add the record, and click **Other New Records.**
3. On the **Resource Record Type** tab, select **Service Location,** and then click **Create Record.**
4. On the **New Resource Record** tab, in **Service,** type **_rvp**. In **Protocol,** select **_tcp**. In **Priority,** type **0**; in **Weight,** type **0**; in **Port number,** type **80**.
5. In **Host offering this service,** type the FQDN. This will be the FQDN of the Instant Messaging router servers in the topology. Click **OK.**

The finished registration will resemble the following example:

```
_rvp Service Location [0][0][80] im.contoso.tld
```

Client Name Resolution

You might decide to prevent internal clients from resolving external DNS entries for security reasons. If so, an lnstant Messaging client will fail on the DNS query for a new contact. For example, a user who wants to add a contact for a particular external domain would be able to resolve a SRV record if the client is able to resolve external DNS entries. However, if the client is blocked from resolving the address, it does not mean that they are unable to conduct an Instant Messaging conversation with the contact from the other organization.

HTTP (port 80) can be made available through the firewalls, even though DNS queries are not allowed. If a client is unable to query a DNS entry for a particular Instant Messaging domain, the user can enter the fully qualified Instant Messaging address for the user (for example, suzan@im.contoso.tld instead of suzan@contoso.tld). In this way the client can be added and conversations can be started even though the client will fail on the initial SRV lookup. Note that when the client has resolved server names, the Instant Messaging Router server is not involved in communications until a new session is initiated. Server and client names are cached for the duration of the session.

Instant Messaging servers use DNS queries to resolve address information for destination connections. If you want to allow communications between internal and external Instant Messaging servers, your company's Instant Messaging servers must always be able to resolve DNS queries for external hosts.

Companies that want to communicate by means of Instant Messaging can do so by configuring routers to work with firewall security and exposing their Instant Messaging infrastructures by using Instant Messaging routers. Clients from one company can connect to servers in other companies by using HTTP, just as they would for internal Instant Messaging. Users can subscribe to presence information when they obtain resolved user names from directories or from aliases provided by correspondents. To configure the firewall settings to allow communication between companies, open System Manager, click **Global Settings,** and right-click **Instant Messaging.**

Scalability for Instant Messaging Servers

Independent of geography or multiple e-mail domain issues, the number of servers increases in direct proportion to the size of the user population. In typical corporate settings, most users will be online concurrently at peak times. In hosting scenarios, only a small fraction, of all users will be concurrently online at any given time.

Therefore, a company with 30,000 active users might need three home servers, because almost all 30,000 users could be logged on at the same time. An ISP with 1 million users and a 5 percent peak online rate would need five home servers, because the peak load would be 50,000 concurrent online users.

Server Location and Network Connectivity

Instant Messaging assumes that home servers and Instant Messaging routers are continuously connected to a WAN and that clients are continuously connected to the network (through low-bandwidth or high-bandwidth connections) while logged on. In the simplest case, all Instant Messaging servers could be placed in one location even if the organization contains subnets in multiple branch offices scattered throughout the region or continent. Clients periodically connect with home servers and Instant Messaging routers, but such connections are low-bandwidth, and of short-duration, and can be made easily over long distances, such as across continents.

Low Bandwidth Connections

If a company has multiple concentrations of users separated by low-bandwidth, or expensive links, you should consider distributing home servers to the highest density areas in the network topology.

For example, if a company has significant populations of users in Asia (3,000 users), South America (3,000 users), North America (11,000 users), and Europe (3,000 users), and if the continents are connected by low-bandwidth or expensive links, the company can deploy two or more home servers in North America and at least one user home server in Asia, South America, and Europe, even though the total number of users worldwide would typically require no more than two home servers.

If a geographically distributed company wants to use Instant Messaging to communicate externally and has only one connection to the Internet, all of the Instant Messaging routers (if more than one has been deployed) should be located near the Internet connection. Instant Messaging clients typically communicate far less frequently with Instant Messaging routers than they do with home servers. Instant Messaging clients cache home server information provided by routers. Instant Messaging routers can be placed at any location that has an Internet connection.

Branch Offices

Special considerations should be made when deploying Instant Messaging in companies with many small branch offices and few (10 to 100) users at each location. Most branch office architectures have a continuous connection to the corporate network backbone, although this connection may be low-bandwidth (56 KB or 128 KB). In such scenarios, branch offices have their own Instant Messaging servers or have multiple servers local to each location.

Clients communicate using the RVP protocol, which is asynchronous and requires low bandwidth, so network traffic is a minor consideration when you determine the number of Instant Messaging servers to deploy and how to position them on your network. The increased traffic caused by adding Instant Messaging to the infrastructure will be minimal in most cases.

You must consider administrative issues, such as the management of servers at remote locations, when you decide where to place your servers. For example, you might have branch offices administer their own Exchange servers. If so, the home server for each location could still be located in the enterprise hub. However. if local administrators want to control collaborative services, Instant Messaging servers can be placed at branch offices. When planning server placement, consider deploying Instant Messaging servers based on the concentration of groups of users that will communicate with one another. You can reduce the number of servers to which a client must subscribe and decrease the associated WAN bandwidth for presence information updates by placing servers within these groups and setting up members of highly communicative groups on single or multiple co-located servers.

Network Address Translation and Internal Firewalls

Instant Messaging clients use IP addresses during connected sessions to initiate requests to servers and to conduct communications with other Instant Messaging clients. Instant Messaging servers query DNS to resolve addresses for server-to-server communications. If a client is on a network that is protected by a firewall that is translating IP addresses, the client will be unable to resolve IP information for clients on another subnet. To resolve this issue, you can place a home server on the subnet protected by network address translation. This allows the home server to act as a relay host for client communications with other clients beyond the local subnet.

Security Considerations

In Instant Messaging conversations, message5 are sent between all parties as Extensible Markup Language (XML)-tagged Multipurpose Internet Mail Extensions (MIME) message body parts over HTTP. Secure MIME (S/MIME) conversations are currently not supported. Neither is SSL or TLS-encrypted sessions with Instant Messaging servers or between Instant Messaging clients. Instant Messaging can be a security issue when users communicate sensitive information in clear text. You can encapsulate an entire communications stream by using a VPN-protected session (such as a PPTP, L2TP, or IPSec tunnel (L2TP and IPSec are available in Windows 2000), but the communications stream might be susceptible to connection issues if the organization is translating network addresses for servers and clients.

Instant Messaging Deployment Recommendations

Further recommendations for deploying lnstant Messaging Service include:

- Keep naming standards for clients, servers, and Instant Messaging domains as simple as possible. Be aware of the use of internal names as described earlier in this chapter. In most cases, you can add the prefix "im" to the SMTP address of users (Sandra@im.contoso.tld) and you should verify that the proper SRV records have been entered in all internal and external DNS servers, if your company wants to use Instant Messaging across a public network.
- Follow SMTP naming conventions in the construction of Instant Messaging names when your company's multiple geographic locations are served by multiple SMTP domains. For example, a U.S. location that has an SMTP domain of us.contoso.tld associated with it should have an Instant Messaging domain of im.us.contoso.tld.
- Deploy one or more IIS servers per Windows 2000 site. These servers can host the required virtual server instances for Instant Messaging router or home servers.
- Deploy one Instant Messaging home server per 10,000 concurrent users, and one Instant Messaging router per 50,000 concurrent users in all but widely geographically distributed companies. In general, the number of home servers increases relative to the number of concurrent users (for centralized deployments). At a minimum, a server hosting this many concurrent user connections requires a dual processor Pentium III or compatible 400-MHz computer with 256 megabytes (MB) of system RAM.

CHAPTER 20

Inter-Organization Replication and Directory Synchronization

Greg Dodge, Senior Consultant, Microsoft

Microsoft Exchange 2000 Server and the Microsoft Windows 2000 Active Directory directory service are extremely flexible when operating in a single Windows 2000 forest. When business requirements demand multiple forest configurations, additional tools and processes are necessary to fulfill the replication requirements. An inter-organization replication requirement exists when two or more Exchange 2000 organizations must be synchronized with each other. Keep in mind that the boundary of an Exchange 2000 organization is the Windows 2000 forest (just as an instance of Active Directory is bounded by a Windows 2000 forest); only one Exchange organization can exist within a forest.

This chapter describes scenarios that require inter-organization replication. This chapter also discusses situations that warrant a directory solution based upon requirements more complex than those for standard synchronization or replication. There may also be situations where the best solution is directory synchronization with a directory from another system. After reading this chapter, you can identify when to use an inter-organization replication solution instead of a synchronization solution, and how to best implement the solutions in your environment.

In This Chapter

Scenarios

Exchange 5.5 Inter-Organization Solutions

Exchange 2000 Inter-Organization Solutions

Microsoft Metadirectory Services

Conclusion

Scenarios

There are two important terms in inter-organization scenarios: *replication* and *synchronization.* Replication is the process of updating objects in a common directory namespace to one or more directory instances. Synchronization is a process used to copy objects from one directory to another in separate directories, with distinct namespaces. In this chapter, the term replication refers to the process that keeps objects updated within the Exchange 5.5 directory or within Active Directory. The term synchronization refers to processes that copy objects between Exchange and another system, such as a Lotus Notes directory or another Exchange directory.

There are different inter-organization synchronization solutions available for Exchange 5.5. Although many of these solutions may work with Exchange Server version 4.0 or Exchange Server version 5.0, this chapter covers only Exchange Server version 5.5 synchronization and replication with Exchange 2000. In Exchange 5.5, each organization has its own directory service. Exchange 2000 is fully integrated with Windows 2000 and uses Active Directory to generate a global address list.

Exchange 5.5 Inter-Organization Solutions

When inter-organization directory synchronization is required in an Exchange 5.5 environment, there are a number of solutions. One might think that there is no reason to do inter-organization synchronization between Exchange organizations because the built-in replication can be used. However, due to limitation? in Exchange 5.5, the replication mechanism can only replicate directories with an Exchange system that has the same organization name. This restriction poses problems when companies merge or when business units decide to connect separate systems.

There are a number of options for these types of situations, ranging from a complete reinstallation of Exchange, to the use of inter-organization synchronization tools. You can use Microsoft Exchange Move Server Wizard to rename one of the systems with the name of the other system, but this process has limitations, especially when more than two systems are involved. This section briefly outlines some of the most commonly used Exchange 5.5 solutions and explains why some cannot be used to do inter-organization synchronization in an Exchange 2000 environment.

This section covers inter-organization synchronization of directory objects such as mailboxes and distribution lists. Some organizations may also need to replicate information in public folders between Exchange organizations, such as the Free/Busy system folder that contains calendaring information. Regardless of which inter-organization solution you choose for Exchange 5.5, the public folders can be replicated using the inter-organization public folder replication service in Exchange 5.5, Service Pack 2 (SP2). This tool uses remote procedure calls (RPCs) to replicate the information in public folders by using a publisher and subscriber model. For more information about inter-organization public folder replication solution, see the Exchange 5.5 SP2 documentation.

Network News Transport Protocol (NNTP) can also be used to replicate folders between organizations. The benefit is that replication can work over the Internet; however. there are security risks.

InterOrg Synchronization Tool

Many companies that routinely purchase and sell entire companies have unique directory requirements. To meet these requirements, Microsoft Consulting Services developed the InterOrg Synchronization tool to synchronize different Exchange 5.5 organizations into a cohesive directory. This InterOrg Synchronization tool is included in the *Microsoft BackOffice* ***Resource Kit, Second edition.***

For servers that are running Microsoft Windows NT Server version 4.0, information about using the InterOrg Synchronization tool is available on the TechNet Web site at http://www.microsoft.com/TechNet. If you have Windows 2000 installed, Active Directory Connector (ADC) helps synchronize multiple Exchange 2000 organizations and Exchange 5.5 sites into a cohesive directory.

> **Note** You cannot use the InterOrg Synchronization tool to synchronize different Exchange 2000 organizations because the tool uses Directory application programming interface (DAPI) technology that will not be supported in the future.

Windows 2000 with Active Directory Connector

Active Directory Connector (ADC) synchronizes Windows 2000 Active Directory with the Microsoft Exchange Server 5.5 directory. This synchronization can be used to aid in the implementation of Active Directory for companies that have deployed Exchange Server 5.5 and is a necessary stage for achieving coexistence between Exchange Server 5.5 and Exchange 2000 Server. ADC is also a stand-alone tool that synchronizes directory entries between two or more Exchange 5.5 organizations in an inter-organization scenario after Exchange 2000 is installed.

ADC is an additional component for Microsoft Windows 2000 Server and Exchange 2000 Server. After installation, a new Windows 2000 service called Microsoft Active Directory Connector appears in Microsoft Management Console. This service can be started and stopped just like any other service. A snap-in and an .msc file, which configures the connection agreements between Active Directory and the Exchange target directory, are also installed.

ADC does the following:

- Uses the Lightweight Directory Access Protocol (LDAP) APIs to perform fast replication between the Exchange 5.5 directory service and Active Directory.
- Hosts all active replication components on Active Directory, not another system.
- Replicates changes only between the directories, whenever possible.
- Maintains object fidelity through replication (for example, the Active Directory Group object matches to the Exchange Distribution List object).
- Hosts multiple connections on a single server and manages these through connection agreements.

Configuring ADC for Inter-Organization Synchronization

Because ADC can replicate Exchange mailbox objects, custom recipient objects, and distribution list objects to Active Directory in an organizational unit, an inter-organization replication solution is easy to build. If you use Active Directory as the intermediary, multiple ADC connection agreements can be set up to replicate organizations, Exchange sites, or recipient containers into an organizational unit in Active Directory. By using separate connection agreements, you can easily replicate those entries to a target recipients container in separate Exchange organizations.

The first step in creating an inter-organization solution for Exchange 5.5 is to install ADC or Exchange 2000 into your production Active Directory forest, then establish an intra-organization connection agreement to the existing Exchange 5.5 environment. Installing Exchange 2000 using the same organization name as the Exchange 5.5 environment requires an intra-organization connection agreement.

You cannot create an inter-organization solution without installing ADC in your production Active Directory forest. For more information about installing ADC and configuring connection agreements within an organization, see "Active Directory Integration and Replication" in this book.

After the intra-organization connection agreement is configured, install as many inter-organization connection agreements as are needed to completely separate Exchange 5.5 organizations. To create an inter-organization connection agreement, click **This is an inter-organizationconnection agreement** on the **Advanced** tab. This tells ADC and Exchange 2000 that the target mailboxes for this connection are in a different Exchange organization. Deploy two one-way connection agreements to create this inter-organization configuration to give you more control over source and target locations for mailboxes.

Caution If you deploy ADC in your production Windows 2000 forest and accidentally configure an inter-organization connection agreement to your Exchange *5.5* organization first, you cannot deploy Exchange 2000 in that forest or create an intra-organization solution with that Exchange *5.5* organization. You can re-create the forest before installing Exchange 2000, or deploy Exchange 2000 in a separate forest. Therefore, focus your migration project on integrating Exchange *5.5* with Exchange 2000 and Active Directory and then deploy inter-organization connections to the other Exchange *5.5* organizations.

Versions *of* ADC

You can synchronize Exchange *5.5* and Active Directory by using the version of ADC that is included with Windows 2000 Server. This version is sufficient for preparing Active Directory to install Exchange 2000 Server. However, when you install Exchange 2000, an updated version is installed. This updated version of ADC keeps Exchange *5.5* and Exchange 2000 synchronized, and has new features to establish the required intra-organization replication between an existing Exchange 5.5 organization and Exchange 2000. The Exchange 2000 version of ADC is also the only version that can be used to provide inter-organization coexistence between multiple Exchange *5.5* organizations and Exchange 2000. The Windows 2000 version does not support inter-organization connection agreements and is not supported by Microsoft if you use it in this manner. Because of limitations in the version of ADC that is included with Exchange 2000, this version cannot provide inter-forest replication between two different Exchange 2000 organizations. However, solutions for this situation are discussed in the following sections.

Exchange 2000 Inter-Organization Solutions

There are a number of options available when you need to establish inter-organization replication between multiple Exchange *5.5* organizations or between Exchange *5.5* and Exchange 2000. When establishing inter-organization replication between two Exchange 2000 organizations, there are a number of challenges that you must overcome.

First, Exchange 2000 does not have its own directory like Exchange *5.5*, and it is dependent on Active Directory. Second, Active Directory does not have built-in functionality that allows replication or synchronization of directory objects between separate forests. Because only a single Exchange 2000 organization can exist in an Active Directory forest, you will need an inter-organization synchronization solution if your company has more than one forest (namespace).

The main reason that only one Exchange 2000 organization can exist per forest is that the schema of a global catalog server must be the satne for all global catalog servers in the entire forest to provide a consistent set of attributes. Because anyone with Schema Admin rights can extend the schema to include a new attribute in the global catalog, replication is restricted within a single forest.

To resolve the issue of replicating directory entries between forests, a utility that is flexible enough to handle different schemas is needed. The utility must also use LDAP or Active Directory Service Interfaces (ADSI) to replicate the entries between the different forests. The following section discusses Microsoft Metadirectory Services, a tool specifically designed to handle inter-forest synchronization requirements.

Microsoft Metadirectory Services

A metadirectory solution differs from the previous replication or synchronization solutions in that it maintains an independent join of objects from one or more directories. In this section, the metadirectory is Active Directory.

Information about people, applications, and resources is scattered throughout most companies. An increasing amount of this identity data is stored in standards-based directory services, but the majority of it remains stored in databases and other specialized forms. Administrators confront identity data management challenges of varying kinds:

- **Global address book applications** These applications typically attempt to synchronize mailbox information between different e-mail directories within a company.
- **Human resources solutions** These solutions quickly propagate information about newly hired employees to all systems that require identity data and also retract the same information when employees leave.
- **E-commerce applications** These applications must synchronize information such as digital certificates for suppliers and extranet users with e-commerce directories outside of company IT firewalls.

With each additional application and platform that you deploy, the number of places where you must manage data increases, which forces companies to manage data in many different locations.

The simplest metadirectory solution is a single directory that holds all information about users, computers, networks, and applications in a company. However, this solution cannot be reached quickly or at any time in at most companies. Therefore, companies seek solutions that link different directory services and applications and provide a consistent way to store, access, and manage their data.

If data must exist in many places, solutions need to provide:

- Connectivity that enables data sharing between multiple directory services, databases, and applications.
- Brokering functionality that distributes changes made in one directory or application to other data repositories in the company affected by the change.
- Integrity mechanisms that maintain owner information, that adhere to integrity rules, and that ensure that related data remains consistent throughout the company.

Because data management issues are so varied, no single solution can resolve all of them. Instead, consider deploying several of the following options:

- Multi-directory access technologies that provide developers with a single programming interface to multiple directory services and database technologies.
- Synchronization connectors that simplify management by keeping pairs of directory services synchronized with each other automatically.
- Directory consolidation strategies that enable companies, over time, to reduce the total number of directories that they need to manage.
- Metadirectory technology that provides companies with advanced connectivity, brokering, and data-integrity management capabilities.

Microsoft Metadirectory Services is the Microsoft metadirectory solution.

You can benefit from a metadirectory by integrating a company's data repositories and joining data shared among them; information about a particular person can then be shared. This is valuable when data from one organization is to be shared with another, where a mail-enabled user in one organization needs to be created as a contact in the other, and where groups in one organization must be represented in another.

Microsoft Metadirectory Services uses management agents to connect to connected directories. By using an integrated discovery and integration process called ***join,*** connector entries are created in the management agent's connector space and joined to a single entry in the *metaverse*. Microsoft Metadirectory Services defines metaverse as that portion of the directory that presents the integrated view of joined objects from multiple connected directories. The metaverse entry can be joined to multiple connected directory entries through these connectors — each held in its respective management agent's connector space. After attributes are joined, they can flow between connected directories and the metadirectory. Microsoft Metadirectory Services supports bi-directional attribute flow, based on a set of rules configured by the administrator. Figure 20.1 depicts an overview of the Microsoft Metadirectory Services structure.

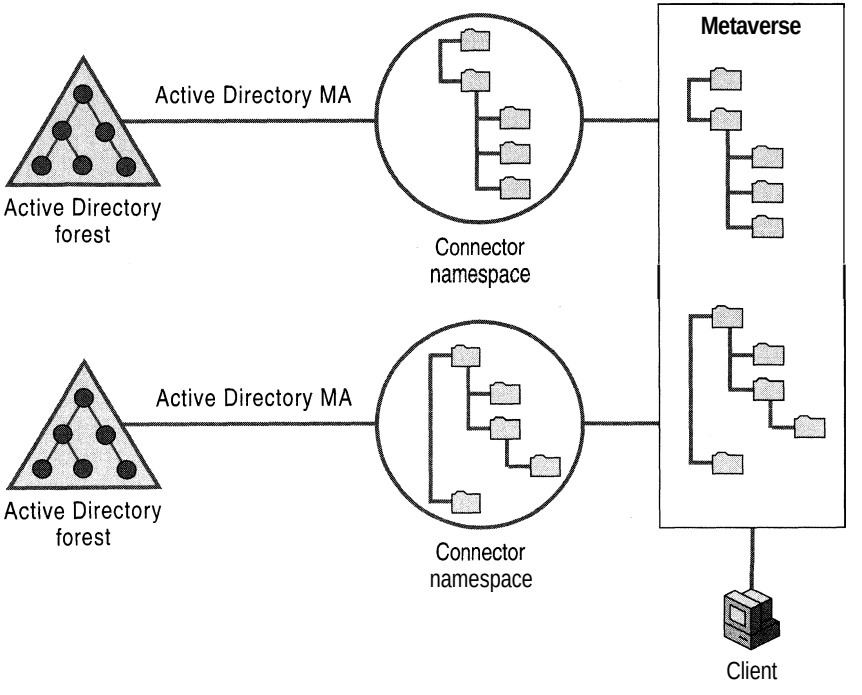


Figure 20.1 Microsoft Metadirectory Services architecture

Microsoft Metadirectory Services provides the infrastructure for customers to integrate directories, messaging systems, network operating systems, databases and resource planning applications. Microsoft Metadirectory Services works with Active Directory to provide a scalable and distributed framework for integrating and synchronizing information between Active Directory and heterogeneous data sources, such as:

- Microsoft Exchange
- Lotus Notes
- Novell Network Directory Services (NDS)
- Microsoft SQL Server
- XML data sources
- Service Advertising Protocol (SAP) systems
- PeopleSoft systems
- Baansystems

A management agent tool kit is fully integrated with Microsoft Metadirectory Services to allow configuration, extension, and creation of new management agents. Microsoft and partners provide consulting services to assist customers deploying and customizing Microsoft Metadirectory Services.

Active Directory Management Agent

Active Directory uses LDAP as its directory access method. Although Microsoft Metadirectory Services supports a generic LDAP management agent for Novell NDS. Netscape, Microsoft Exchange 5.5, Isocor and other connected directories, Microsoft Metadirectory Services is not optimized for Active Directory. The generic LDAP management agent needs to read the entire directory content to infer changes that occurred in an LDAP-based directory. This approach is inefficient and is not always practical in large directory environments. This is particularly true when administrative limits are used. Administrative limits restrict the number of entries returned by a single search. The common practice of setting these administrative limits requires the LDAP management agent to automatically execute many smaller queries as part of a procedure to read the entire content of the directory.

Microsoft plans to provide the Active Directory management agent, a management agent that is fully integrated with and optimized for Active Directory, with Microsoft Metadirectory Services version 2.2. The Active Directory management agent uses the Microsoft LDAP directory synchronization control to obtain directory updates from Active Directory, and uses LDAP to flow attributes and create objects in Active Directory. This control allows the Active Directory management agent to request the changes that have occurred in Active Directory without reading the entire content of the directory. This new mechanism allows *near real-time* synchronization with Active Directory because it is now possible to schedule synchronization updates as often as every minute. Unlike Microsoft Metadirectory Services 2.1 management agents, the Active Directory management agent interacts directly with the synchronization engine and automatically detects schema changes that occur in Active Directory. The architecture of the Active Directory management agent allows it to automatically adapt to schema changes without requiring an adaptation of the management agent. In addition, greater efficiency is realized because the Active Directory management agent uses no intermediate import and export file, as is the case with file-based management agents.

The Active Directory management agent in Microsoft Metadirectory Services 2.2 can be used to directly mirror the contents of a Windows 2000 forest. When the Active Directory management agent mirrors multiple trees of a forest or multiple forests, it is best to create a virtual metaverse root for each by using the dc=com namespace as the parent for each tree or forest. This allows each tree or forest to be fully replicated to Microsoft Metadirectory Services, preserving its namespace while still allowing for replication between trees or forests by Microsoft Meladirectory Services. For more information about creating your namespace when working with multiple trees or forests, see the Microsoft Metadirectory Services documentation.

Inter-Forest Scenarios

There are several reasons why you might need to limit transitive trust and schema proliferation, such as mergers, acquisitions, joint ventures, and so on. This need will require you to employ multiple Windows 2000 forests. The Active Directory management agent can create a single namespace for these separate forests by integrating information from multiple line of business applications into Active Directory. The Active Directory management agent integrates and synchronizes objects and attributes between multiple forests.

Multiple forest environments are complex. Many Active Directory features that are supported in single forest deployments are not supported in multiple forest configurations. For example, shared schema and security principals are not supported in multiple forests. However, you can reference another forest's security principals on access control lists (ACLs), as long as trust is established. Although establishing trust may be an issue in many inter-forest scenarios, it can be a useful solution to inter-organization collaboration problems.

Support for the following two multiple-forest company scenarios is planned for the Microsoft Metadirectory Services 2.2 release of the Active Directory management agent:

- Centrally managed Active Directory forests
- Peer forests

Centrally Managed Forests

The following scenario lists most inter-forest Active Directory requirements. The goal is to integrate and manage multiple Active Directory forests from an extended entity (one company).

The company in question needs to:

- Manage identities (users and contacts) across forests, which includes creating and deleting user accounts when an employee joins or leaves the company.
- Synchronize e-mail address book attributes between the messaging systems of each forest (Exchange 2000 in this case).
- Manage distribution groups across forests.
- Synchronize inter-forest locator information, such as site and subnet directory entries, to support printer identification and searching.

In this scenario, the Active Directory management agent and Microsoft Metadirectory Services manage various entities: user, contact, organizational unit, group, and public folder objects in Active Directory. The metadirectory manages entries that originate from a central human resources system in the following way: when the human resources department adds an employee (user), business rules determine in which forest domain an account is created. Users are added to appropriate domains and corresponding contacts are added to a domain in other forests.

Not all objects in a forest need to be propagated to other forests: only a subset is required. This is achieved by identifying individual containers of interest, and types of objects to propagate from one forest to the other.

Match objects from one forest to the other as follows:

- Users of the source forest will be projected as contacts in the target forest; they can also be projected as disabled users.
- Contacts of the source forest will be projected as contacts in the target forest.
- Security and distribution groups of a forest will be projected as universal distribution groups in the target forest.
- Organizational units of the source forest will be projected as organizational units in the target forest.
- Domain controllers of the source forest will be projected as organizational units in the target forest.

Some organizations may also need to support site, subnet, and location information and object synchronization between forests to support inter-forest resource locator services. The first to make use of this information will be the printer locator service. This synchronization significantly reduces the time required to deploy Windows 2000 across multiple forests. After this information is available, users can locate the closest printer when using the Add Printer wizard. Subnet and location information needs to be projected in the root domain configuration container of each forest. After this information is available, users can locate printers in a uniform and consistent manner regardless of which forest or physical location they use to log on to Windows 2000.

Peer Forests

There are occasions when multiple forests are peers and need to be integrated. Each forest is individually managed; therefore, these forests are peers and not managed from a central, external source.

For Exchange 2000, this implies that as users and groups are added to one organization, they need to be projected and their attributes need to be synchronized with the other organizations.

The forests have a many-to-many relationship instead of a one-to-many relationship. Aside from migrating information that ADC creates, enter only a few parameters with little customization to make this scenario work with the Active Directory management agent. Create one management agent per forest; identify what object types to project, where to project them, and how often to perform the synchronization. This process is described in the following section.

Setting Up Microsoft Metadirectory Services

The example for this section is that two peer Windows 2000 forests must be integrated after a merger. In this situation, two or more Exchange 2000 forests need to be synchronized to support messaging between users in the two forests.

To create an inter-organizational replication solution with Microsoft Metadirectory Services 2.2 you need to configure one Active Directory management agent per forest. Configure this management agent in reflector mode to ensure that Windows 2000 users are projected to the metaverse. For more information about Microsoft Metadirectory Services and its management agents and about configuration of the Active Directory management agent, see the Microsoft Metadirectory Services documentation.

The following are the requirements for the peer forest scenario:

- Directory synchronization between each forest must occur.
- Contact entries created by Microsoft Metadirectory Services in either Exchange organization must mirror the corresponding entry, attribute for attribute.
- Ownership of all Exchange entry attributes must remain with the original Windows 2000 forest.

Designing Microsoft Metadirectory Services

You can install Microsoft Metadirectory Services on a server in the target forest. Two Active Directory management agents have been created and configured to communicate with the different forests.

The management agent for each forest brings all mailboxes and lists into the metaverse. It also extracts the other forest's entries from the metaverse and imports them. This management agent is configured to point to a specific Active Directory server.

By default, all pertinent Exchange attributes are populated to the metaverse. With the attribute-flow user interface, you can customize additional attributes to populate the metaverse. The default configuration populates metaverse entry attributes to the contacts and groups of the target forest.

Create resource objects for each Active Directory forest in which you want to create contacts and groups This resource dictates into which Active Directory forest, Active Directory domain, and Active Directory container these objects project and synchronize. There is an option to create a flat list of ucers and groups, or to inherit the hierarchical structure of the source forest.

Assign these resources to subtrees in the metaverse or, by using Microsoft Metadirectory Services scripting, assign them to specific entries.

- **Option 1** Select a subtree in the metaverse (which represents an Active Directory forest and Exchange organization) and assign it a resource. This causes the creation of all objects of interest (contacts, groups, sites, subnets) in this forest and the synchronization of object attributes in the container of the target forest specified by the resource.

  By setting up a single resource per forest and assigning each resource to subtrees of the metaverse representing the forests, you cause Microsoft Metadirectory Services to create these objects in the target forest. Attribute flow then takes over and ensures that the attributes are kept in synchronization.

- **Option 2** You can script specific behavior based on attributes of an object. For example, you might only want to synchronize objects of one forest to others based on a particular attribute or a particular value. This might imply that only full-time employees can be synchronized with a partner's Exchange 2000 organization. See the Microsoft Metadirectory Services documentation to determine how to use these features.

To automate directory synchronization between two Windows 2000 forests, the management agents need to operate in delta-operation mode on a regular basis. Microsoft Metadirectory Services includes a scheduling component that can schedule the management agents.

Conclusion

The inter-organization tools that work well depend on Exchange 5.5 DAPI or work only with the Exchange 5.5 Directory Service. Windows 2000 and Exchange 2000 include ADC that uses LDAP to synchronize between Exchange 5.5 and Active Directory, but ADC currently does not support inter-forest synchronization.

Microsoft Metadirectory Services can combine the mailbox-enabled users and groups that originate in each Exchange organization into a single directory (The Microsoft Metadirectory Services metaverse) while keeping track of which organization owns each directory entry. Microsoft Metadirectory Services can update any of the connected forests with entries outside of each Windows 2000 forest namespace. Through attribute flow, administrators have control over which attributes synchronize between the forests to support the needs of each company. In addition to these features, Microsoft Metadirectory Services also provides advanced tools that allow more complex business rules to support provisioning, employee hiring, and termination. For more information about implementing Microsoft Metadirectory Services, search the Microsoft Web site at http://www.microsoft.com.

CHAPTER 21

Branch Office Scenarios

Mark Garcia, Consultant II, Microsoft
Dan Bloch, Senior Consultant, Microsoft

A number of issues arise when you deploy Microsoft Exchange 2000 Server in branch offices. These issues include organization of Microsoft Windows 2000 subnets, Active Directory replication, Exchange 2000 routing group design, and the placement of Domain Name System (DNS) servers, domain controllers, and global catalog servers. You can choose from several administrative approaches when you deploy Exchange in branch offices, including a centralized, decentralized, or hybrid approach.

This chapter addresses these issues and discusses Exchange 2000 and Windows 2000 administrative options. It also discusses the available clients and assesses them in the context of branch office deployments.

In This Chapter

Replication and Routing Group Dependencies in Windows 2000

Branch Office Administrative Models

Global Catalog Placement

Branch Office Messaging Client Considerations

Replication and Routing Group Dependencies in windows 2000

Because the global address book for Exchange now resider in Active Directory, you must understand how Windows 2000 site architecture affects Exchange 2000 deployments.

A Windows 2000 site consists of a set of Internet Protocol (IP) subnets with fast, reliable connectivity. (Networks with LAN speed or greater are considered fast networks.) Windows 2000 sites are used to control Active Directory replication across slow WAN links and to locate the nearest domain controllers. Active Directory clients and servers use the forest site topology to route query and replication traffic efficiently. Forests provide routing and replication benefits because the collection of one or more Windows 2000 domains share a common schema, configuration, and global catalog, and because they are linked with two-way transitive trusts. A site topology also helps you decide where to place domain controllers on your network. For example, you should place a greater number of domain controllers in the area of your network in which user load is highest. Exchange 2000 users can use Windows 2000 sites to locate domain controllers for domain authentication and subsequent access to their Exchange mailbox.

Exchange 2000 Routing Groups and Administrative Groups

Previously, when you deployed Exchange 5.5 servers in a branch office, you had to consider which Exchange 5.5 site to place the Exchange 5.5 server in, and which connector would connect different sites, based on available bandwidth. The Exchange 5.5 concept of a site no longer exists in Exchange 2000. In Exchange, you can use administrative groups and routing groups to separate the logical and physical organization of your company. You can use administrative groups to organize the servers according to the administrative structure used in your organization. You can also use routing groups to map out the physical network of connections between servers.

There is no administrative correlation between administrative groups and routing groups. Therefore, in the branch office environment, you could assign administration of Exchange 2000 mailboxes and server configuration to branch office administrators. and assign Exchange 2000 routing architecture to central mesraging backbone administrators. Figure 21.1 illustrates this concept; it shows four administrative groups managed by the individual branch offices and six Routing groups managed centrally by the messaging backbone administrators.

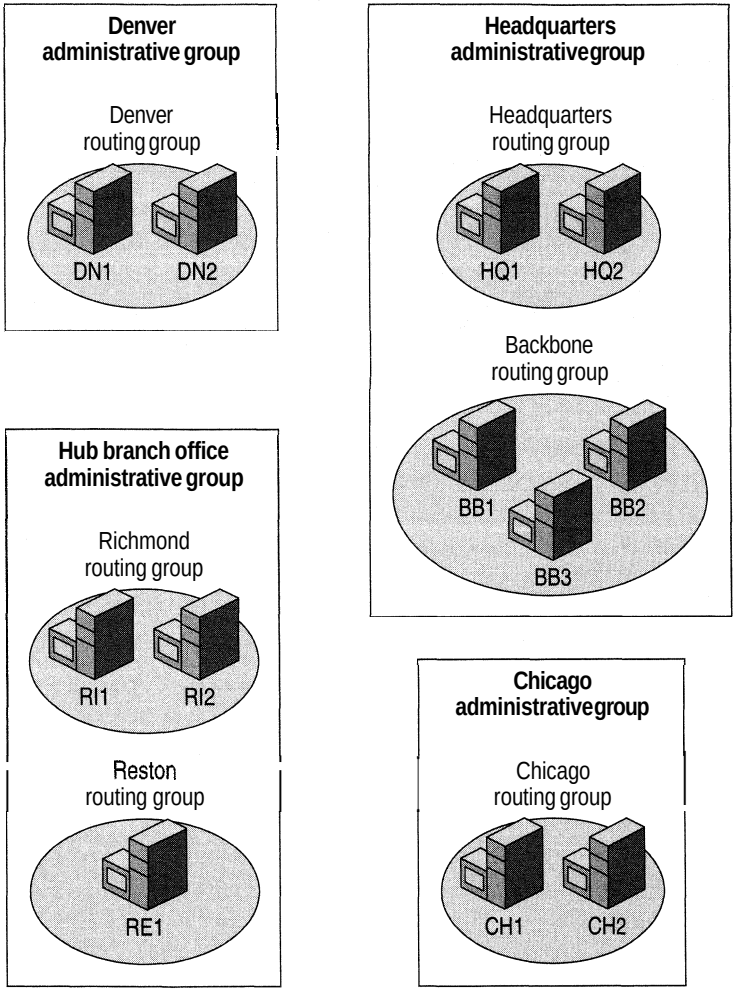


Figure 21.1 Examples of administrative groups and routing groups

Designing Routing Groups

Before you design Exchange 2000 routing groups for the branch office, you must understand a few concepts that are new to Exchange 2000. One of these concepts is that all Exchange 2000 servers in a routing group must have a permanent connection and must be able to contact the designated *routing group master*. The routing group master is a server in the routing group that is responsible for maintaining the link state table of all other routing groups in the Exchange 2000 organization.

Mapping Windows 2000 Sites to Exchange 2000 Routing Groups

Initially, you can base your Exchange 2000 routing group design on your existing Windows 2000 site architecture, but the two designs are likely to diverge at some point. This is because a Windows 2000 site bases communication between servers on a remote procedure call (RPC), whereas an Exchange 2000 routing group bases communication between servers on Simple Mail Transfer Protocol (SMTP). SMTP-based communication requires much less available bandwidth and therefore requires fewer Exchange 2000 routing groups than the number of subnets required on a Windows 2000 site. In addition, Exchange servers in an Exchange 2000 routing group do not require high available bandwidth (for example, 128 Kbps) in the same way as servers on an Exchange 5.5. site. However, Exchange 2000 servers in a routing group should have a persistent high-speed connection between them.

> **Tip** Based on this information, if you have a branch office that has an unreliable network link to other offices, consider making it a separate routing group. Likewise, if you have several offices with high network reliability, you can place them within one routing group.

Bridgehead Routing Group Servers

It is also important to designate routing group bridgehead servers for replication. With Exchange 2000, you can now designate bridgehead servers when connecting two Exchange 2000 routing groups. To build a high level of redundancy, you should designate at least one primary bridgehead server and one secondary bridgehead server per routing group connection. Another option is to allow the routing group master to select the bridgehead server.

Routing Groups in Exchange 2000 Mixed Mode Versus Native Mode

An Exchange 2000 organization can operate in two modes: native mode and mixed mode. Native mode offers full Exchange 2000 functionality, whereas mixed mode offers interoperability between Exchange 2000 and previous vercions of Exchange. When you install Exchange 2000, your Exchange organization is in mixed mode by default. This default setting ensures future interoperability with previous versions of Exchange (for example, the installation of an Exchange 5.5 server at a later time), even if no Exchange 5.5 or earlier servers exist in your organization at the time of installation.

The concept of mixed mode and native mode organizations in Exchange is similar to the concept of mixed mode and native mode domains in Active Directory. With Exchange 2000, you can select native mode or mixed mode only at an organizational level. However, no direct relationship exists between the mode of an Active Directory domain and the mode of an Exchange organization. The similarity exists only in terms of naming and restrictions on earlier versions.

For Exchange Server 5.5 and Exchange 2000 to coexist and replicate directory information, the Exchange 2000 configuration must remain in a state that can be recognized by Exchange Server 5.5. Active Directory Connector (ADC) is critical to ensure coexistence with previous versions of Exchange.

By default, Exchange 2000 is installed in mixed mode. If you are upgrading from Exchange 5.5, when you first install Exchange 2000. there is a one-to-one mapping of Exchange 5.5 sites to Exchange 2000 administrative and routing groups. This means that each Exchange 5.5 site, which could be a branch office, is designated as one administrative and routing group in Exchange 2000. However, moving servers and maiiboxes between these groups is not simple. For more information, see "Moving Exchange Servers" in this book.

> **Note** All Exchange 2000 servers placed in routing groups must be in the same Windows 2000 forest. You can, however, locate Exchange 2000 servers from multiple Windows 2000 domains in a single routing group.

Branch Office Administrative Models

You can choose from the following approaches to administer Exchange 2000 for the branch office:

- Centralized administration from one location
- Distributed administration from multiple branch offices
- A combination of centralized and distributed administration

As you plan your Exchange 2000 administration, keep in mind that Windows 2000 is closely tied to Exchange 2000. Therefore, you must also consider who will be responsible for Windows 2000 account and group administration.

PlanningY Exchang 2000 1 v d

Before you deploy Exchange 2000, you must consider these five planning issues:

- Present administration model
- Network considerations
- Client access methods
- Security
- Active Directory permissions

Present Administration Model

Before choosing which branch office administrative model to use for Exchange 2000, you should identify how the following tasks are accomplished in your organization and who performs each of these tasks:

- Windows NT 4.0 or Windows 2000 account and group administration
- Mailbox administration
- Exchange or messaging server maintenance and configuration
- Site Connector administration or messaging backbone administration

Network Considerations

When designing your Exchange 2000 environment, you should make sure you have an accurate and up-to-date diagram of the network between your branch offices and your central office. Also, you should be aware of any unreliable network connections between branch offices. This information forms the foundation of your Exchange 2000 routing group design.

Different Types of Mail Client Access

It is important to classify the types of mail users you have in each of your branch offices. This is useful in determining the administrative approach, routing group design, and type of mail client most appropriate for each branch. You could, for example, provide light client access to e-mail through Outlook Web Access in a front-end/back-end configuration in which the only requirement is a browser capable of frames and JavaScript or ECMAScript, or you could deploy the full Outlook 2000 client. Identify the approach you are going to take with each branch office prior to deployment. For more information, see "Branch Office Messaging Client Considerations" later in this chapter.

Security

When you consider the various types of e-mail client access methods available with Exchange 2000, it is important to understand which access methods are secure and what additional security options are available. With Outlook Web Access, you should use Secure Sockets Layer (SSL) through Web server certificates to provide secure access from the Internet. With Post Office Protocol version 3 (POP3) or Internet Message Access Protocol version 4 (IMAP4) access from the Internet, you should use a client-side certificate to provide a secure channel. In addition, because the routing group connectors are SMTP-based, you should encrypt your routing group connections through IPSec, which is a set of protocols that support secure exchange of information

at the IP layer. Although IPSec supports two encryption modes — transport and tunnel — transport mode encrypts only the data portion of each packet, leaving the recipient and routing information accessible. You should use the more secure tunnel mode, which encrypts the recipient and routing information, as well as the message content. For IPSec to work, the sending and receiving devices must share a public key. This allows the receiver to obtain a public key and authenticate the sender using digital certificates.

With sufficient planning, based on the above considerations, you should be able to determine a general approach to Exchange 2000 administration. There are, however, several new approaches to administering Exchange 2000 that may change your current administrative model.

Active Directory Permissions

You should understand which Active Directory permissions are appropriate to grant to your administrators for mailbox and account administration. There are several precise directory permissions available for your central administrators. Active Directory in Windows 2000 allows for very flexible mixtures of account and Exchange 2000 mailbox administration, so you should involve different administrative roles early in your design to ensure that you provide adequate permissions. In addition, you can use both the Windows 2000 Delegation Wizard and the Exchange 2000 Delegation Wizard to grant varying levels of Exchange and Windows 2000 administrative permissions.

Centralized Exchange 2000 Administrative Model

With a centralized Exchange 2000 administrative approach for your branch offices, all Exchange 2000 administrative groups and routing groups are managed from a central location.

The centralized administrative approach works best when each branch office is connected to the central office with high network reliability and availability, but this is not a requirement.

Routing Group Strategies

To simplify centralized administration, it is best to keep as many mailbox servers as possible in the central office and then place these mailbox sewers in a central hub routing group. Keep in mind that each server must maintain high network availability between all other servers in this routing group.

In addition to mailbox-only servers, you should place collaboration or application servers, such as dedicated public folder servers, in this hub routing group. Finally, if you cannot place all mailbox servers in this hub routing group and you must create other routing groups, you should designate one or two dedicated connector servers in this hub routing group to be the bridgehead connector server or servers that connect all other branch office routing groups.

Central Front-End/Back-End Architecture and Routing Groups

In this centralized administrative approach, you first determine if there are branch offices in your environment that don't have messaging administration or in which you don't want to place messaging servers. This is important because in a front-endhack-end configuration, a bank of protocol servers handles the incoming client connections (front-end servers) while the store servers are dedicated to running the databases (back-end servers). Front-endlback-end architecture provides a unified namespace so that if you have multiple servers and user mailboxes are replicated to all servers, a user does not have to know the name of each specific server or which one to log on to. For example, if you have three Exchange servers that have a copy of all user mailboxes, normally you divide the user load by configuring certain users to connect to server 1, others to server 2, and the rest to server 3. If all servers are part of a front-endlback-end configuration, a single name provides user access to all three servers in your configuration. This configuration allows software or hardware load balancing to randomly distribute the load among the three servers without having to reconfigure clients. However, if you have clients in the branch office that must log on to a specific server, do not use a front-endhack-end configuration.

If there are branches that don't need specialized administration, and you determine that they don't need full Outlook 2000 client functionality, you can create a central front-endlback-end architecture. The branches connect to the front-end servers over the WAN to retrieve mail and calendar information. It is important to remember that servers in the front-endlback-end configuration work only with Internet protocols such as IMAP, POP, and HTTP, and not the MAPI protocol, which is used by Outlook.

Placing front-end and back-end servers in a second routing group located in the central office provides redundancy. This central routing group for the branch offices contains front-end Exchange servers such as Outlook Web Access, IMAP4, and POP3 servers for branch office users to access over the Internet or intranet, and will also contain back-end application and mailbox servers. Because the front-end servers will perform queries and address lookups, you should configure these servers as global catalog servers to improve response time. The front-endhack-end approach can greatly reduce administrative cost, client deployment time, and maintenance. For an example, see Figure 21.2.

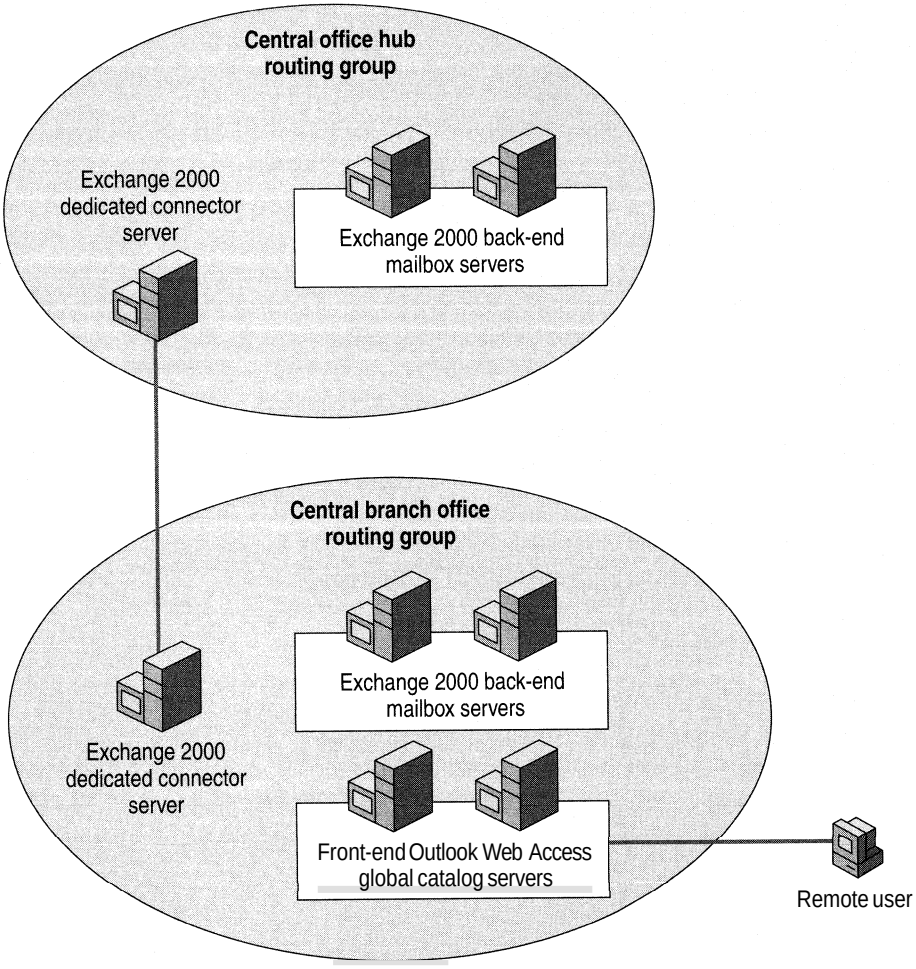


Figure 21.2 Central office routing groups

Administrative Group Strategies

Because your routing group design does not have to be related to your administrative group design, you should simplify your administrative group architecture to make centralized administration easier. Placing all of your Exchange 2000 servers in one administrative group is your best

approach because one administrative group can contain several routing groups. In Exchange 2000 Server, an administrative group is a collection of servers and routing groups that share an administrative security context. All servers or routing groups that you want to be administered by the same group of administrators can be placed in the same administrative group. For example, two servers that support sales staff are located in New York and London; one belongs to the New York routing group and the other to the London routing group. Add both the New York and London

routing groups to the Sales administrative group you would add both the New York and London routing groups so that authorized administrators of the Sales group can make changes to the servers dedicated to the sales staff. You can't move servers between administrative groups so it is important to plan carefully before assigning servers to administrative groups.

Windows 2000 Account Administration and Exchange 2000 Mailbox Administration

If you plan to centralize administration of all Windows 2000 accounts and Exchange 2000 mailboxes, you should create a simple organizational hierarchy that uses only a few layers of organizational units. In other words, avoid nesting your accounts and mailboxes several layers deep in the Active Directory hierarchy. This simplifies account and mailbox administration because the accounts are managed in fewer locations. Consider placing all your Windows 2000 accounts and Exchange 2000 mailboxes in a couple of high-level organizational units. For an example of a simplified hierarchy, see Figure 21.3.

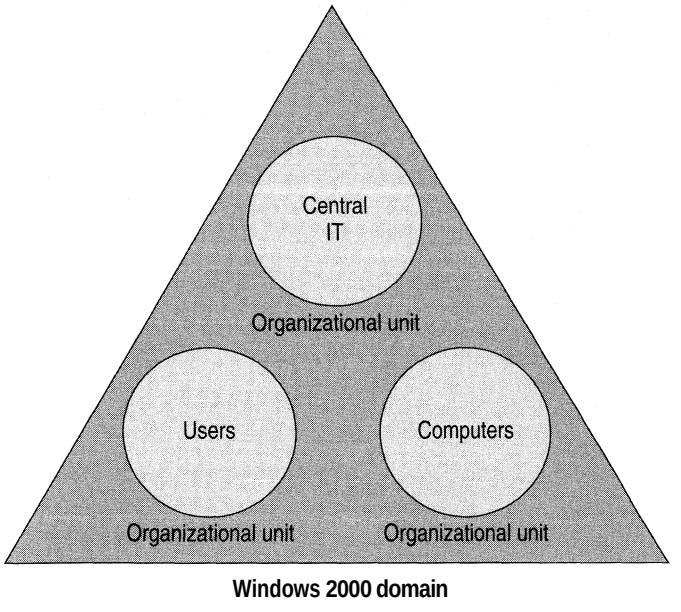


Figure 21.3 A simplified Active Directory organizational unit structure

Distributed Exchange 2000 Administrative Model

In the distributed Exchange 2000 administrative model, all Exchange 2000 administrative groups and routing groups are managed from each branch office or division location.

The distributed administrative approach works best when each branch office or division is administered separately already.

Routing Group Strategies

When using distributed administration with Exchange 2000, you should create routing groups based on how branch offices are connected. Take into account the network reliability between offices; when the link between offices is not reliable, you should create a separate routing group.

Depending on the total number of routing groups, each distributed routing group probably contains mailbox-only servers; it may also contain servers dedicated to routing group connectors. In addition, you can place collaboration or application servers, such as dedicated public folder servers, in the distributed routing group.

> **Tip** When designing routing group connectors, you should always configure at least two paths to every other routing group to provide redundant routing paths in case a network link is not available or a bridgehead server goes offline. Exchange 2000 can re-route messages based on link state information and can automatically redirect messages destined for a link path that is unavailable.

Figure 21.4 illustrates routing group connector redundancy.

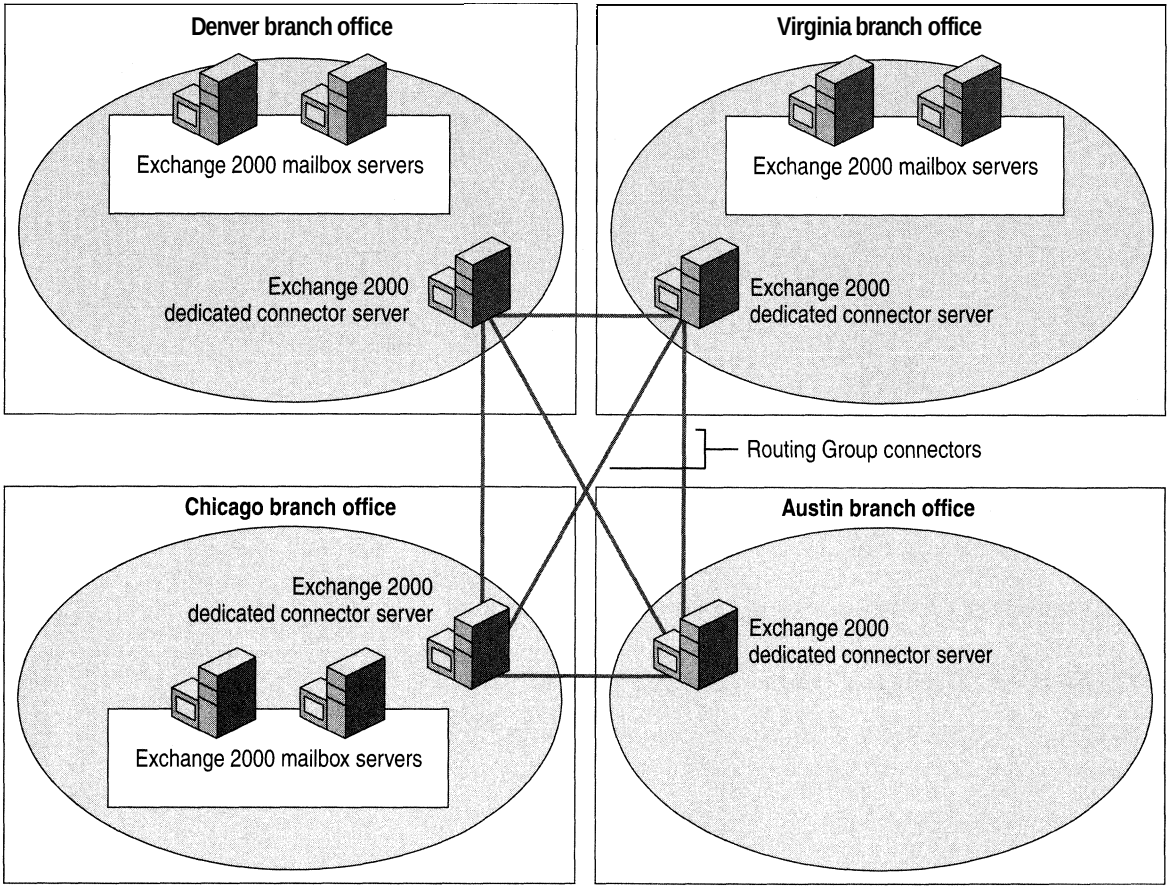


Figure 21.4 Redundant routing group connectors

Each connector server in Figure 21.4 has three paths to another routing group. This ensures that information flows reliably between all branches. For example, if the routing group connectors from the Denver to Virginia and Virginia to Austin branch offices are not available, the routing group connector from Denver to Austin provides another option.

Administrative Group Strategies

In a distributed administrative model, your administrative groups can match your branch offices or divisions. You should place Exchange 2000 servers in each administrative group, based on your regional or branch administration model. As mentioned previously, moving servers between administrative groups is not a simple matter, so plan carefully before assigning servers to administrative groups.

Windows 2000 Account Administration and Exchange 2000 Mailbox Administration Strategies

If you want to distribute Windows 2000 account administration and Exchange 2000 mailbox administration, you should place your accounts and mailboxes in division-based or branch-based organizational units to provide a more precise set of Group Policy and directory permissions based on divisions or branch offices.

Mixed Centralized and Distributed Exchange 2000 Administrative Model

When you choose an administrative approach for your branch office, all Exchange 2000 administrative groups and routing groups can be managed from a combination of central and distributed locations.

The mixed administrative approach works best when each branch office is connected to your central office with a highly reliable network connection, but this is not a requirement. You must define each administrative role in your organization carefully to determine which server components to assign to a central location and which to distribute to remote locations.

Routing Group Strategies

In a mixed administration model, you can choose from several routing group approaches. One approach is to assign as many mailbox servers as possible to the central office, and then place all of these mailbox servers in a central hub routing group. Distribute the remaining routing groups to your branch offices so that each branch can maintain its own messaging infrastructure. When using this approach, keep in mind that each server must maintain high network availability between all other servers in this routing group.

In addition to mailbox-only servers, place collaboration or application servers, such as dedicated public folder servers, in your hub routing group. If you cannot place all mailbox servers into the hub routing group and must create other routing groups due to poor network resiliency, you should assign one or two dedicated connector servers to act as bridgehead servers to all other branch office routing groups.

If your messaging or networking backbone administrators are located centrally, you can also use a mixed approach that assigns all routing groups to a central location.

Central Front-End/Back-End Architecture and Routing Groups

In this mixed administrative approach, you first determine if there are branch offices in your environment that don't have any messaging administration. If so, and you determine these offices don't need full Outlook 2000 client functionality, you can create a central front-endhack-end architecture to process mail and calendar information. Remember, however, that servers in the front-endhack-end configuration work only with Internet protocols such as **IMAP,** POP, and HTTP, and not the MAPI protocol, which is used by Outlook.

Placing these front-endhack-end servers in an alternate routing group located in the central office makes them easier to administer. This central routing group for the branch offices contains front-end Exchange servers such as Outlook Web Access, IMAP4, and POP3 servers for branch office users to access over the Internet or intranet, and also contains back-end application and mailbox servers. Because the front-end servers are performing queries and address lookups, you should configure these servers as global catalog servers to improve response time. The front-endhack-end approach can greatly reduce administrative cost, client deployment time, and maintenance.

Figure 21.5 shows how both the central headquarters hub routing group and front-endback-end routing group, located at the headquarters central office, connect to a mixture of distributed branch office routing groups.

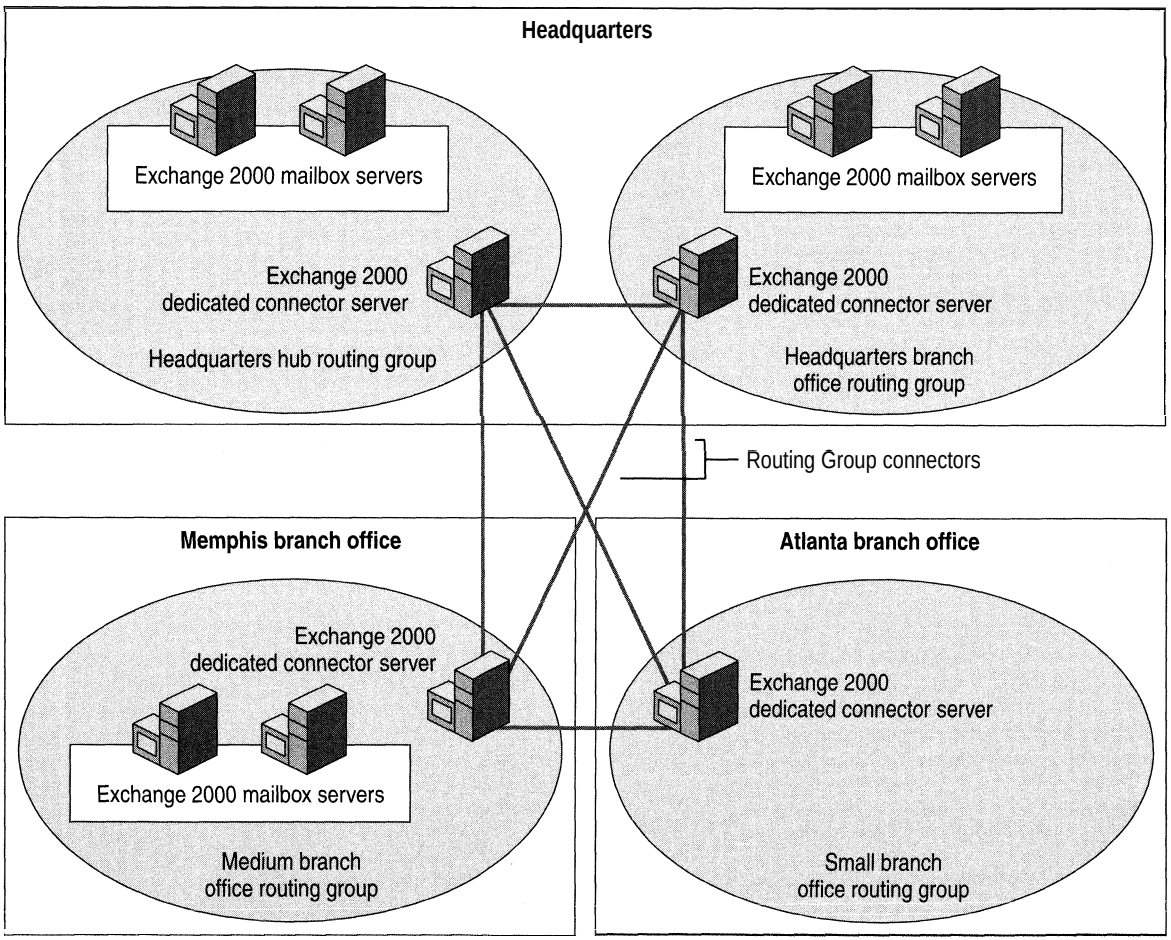


Figure 21.5 Mixed centralized and distributed routing groups

Administrative Groups

With a mixed administrative model, you can take a number of different administrative approaches. You should simplify your administrative group architecture to make overall administration easier. Placing all your Exchange 2000 servers in one administrative group allows servers from multiple routing groups to be grouped together and simplifies management of permissions. For example, after you create an administrative group and set permissions for it, any new objects you add inherit the permissions you set for the group. This can be significant if you have one administrator who manages multiple servers in multiple locations. If you have five Exchange servers, it is easier to define a set of permissions for one administrative group that contains the five servers than it is to define the same set of permissions separately for each server. Because routing groups are used to define the physical network topology only, using one administrative group simplifies the management process.

> **Note** You cannot move servers between administrative groups, so plan carefully before assigning servers to administrative groups. Moving servers between routing groups is an easy process.

Windows 2000 Account Administration and Exchange 2000 Mailbox Administration

If you choose a mixture of centralized and distributed administration, all Windows 2000 account administration and Exchange 2000 mailbox administration is either centralized or distributed. The approach you take determines the placement of your Exchange 2000 mailboxes and Windows 2000 accounts in the Active Directory organizational unit structure. You could place all your accounts and mailboxes in one or more high-level organizational units or in several organizational units based on branch office or division.

Global Catalog Placement

Because Exchange 2000 relies heavily on a global catalog server on both the Outlook 2000 client and server side, you should assign a minimum of one global catalog server for every Windows 2000 site. Two global catalog servers per site provide an even higher degree of redundancy. Depending on available bandwidth between each branch office and hub location, and the number of Exchange client queries, you can place a dedicated domain controller in each location or routing group with a global catalog replica to ensure that Exchange 2000 clients and servers can provide services through the global catalog. In addition, you can install a global catalog server on front-end servers such as Outlook Web Access to help provide faster client response times. You should note that adding a global catalog increases hardware requirements (CPU, RAM, and so forth) for a front-end server. See Figure 21.6 for an example of global catalog server placement in an Exchange 2000 organization.

Close access to a global catalog enables both local mailbox users and local Exchange servers to access the necessary Active Directory information, without requiring RPCs across a potentially slow WAN link.

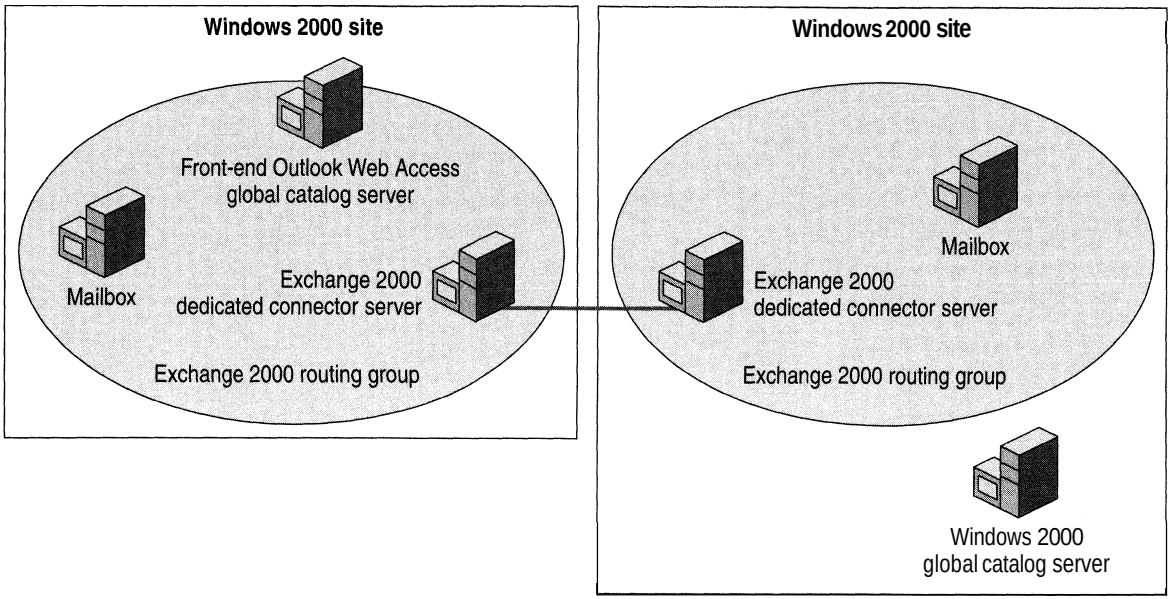


Figure 21.6 Global catalog server placement in branch offices

Branch Office Messaging Client Considerations

Before you select client types for your branch offices, you should determine your deployment objectives and functionality (user) requirements. For example, your company might be concerned about network reliability and cost efficiency. The staff might require availability and format of Internet mail sent to your users. You must consider these various and sometimes competing requirements when selecting the appropriate messaging client type for your branch offices. After determining your requirements, you can set your priorities.

This section presents the methodology for mapping the deployment objectives and user requirements—the design goals—to the messaging client functionality matrix.

Deployment Objectives

You should determine your deployment objectives first. This helps you to determine functionality requirements. Consider the following sample business requirements and then rank them:

- Ability to communicate and share documents with people inside and outside the company
- E-mail addresses that identify business roles rather than a person (no personal names in e-mail addresses)
- Low deployment costs
- Branch office users must be excluded from the central directory
- Central messaging storage

Functionality Requirements

Users often know precisely how they use computers and what they require. Consider the following sample functionality requirements and then rank them:

- Roving and telecommuting capabilities
- Offline capabilities
- Security
- Directory, calendar, and personal information management functionality
- Single interface for all applications

Mobile vs. Offline Requirements

Mobile and offline computer usage requirements are an important consideration. You should understand the difference between mobile and offline users. Mobile users can use more than one device or workstation to connect to mail and will likely want their preferences preserved regardless of their point of entry. Offline users may or may not be mobile users. The difference is that offline users use their mail store without being connected to the network. Each method presents different constraints; however, it is possible to devise a workable solution that permits both forms of access. The following sections discuss various messaging client solutions.

Outlook 2000

Outlook 2000 is the Microsoft Exchange client included with Microsoft Office 2000. It is a full MAPI client that uses protocols such as RPC. It generally assumes a high connectivity network such as a LAN, although it is possible to support users over slower links if you consider the ratio of scalability to available bandwidth. Outlook 2000 takes advantage of the full functionality of Exchange, including access to the public and private information stores, as well as all other client functions, such as calendar information, contacts, tasks, and notes.

Advantages

Advantages of deploying Outlook 2000 include:

- Provides a large feature set such as spell check, search, and rules
- Does not depend on network links for processing because most of the processing occurs on the local workstation, resulting in a faster user experience for locally rendered activities such as composing a message
- Provides access to Notes, Journal, and Sticky Notes (not available in other solutions except Terminal Services)
- Provides the ability to manage folders
- Provides user-friendly interfaces for various functions such as calendar information or contacts
- Provides an excellent offline solution not only for reading and sending messages but also for folder and message management
- Supports Secure Multipurpose Internet Mail Extensions (S/MIME) encryption and signing
- Integrated with Office 2000
- Takes full advantage of the Windows 95, Windows 98, Windows NT, and Windows 2000 operating systems
- Easily configured in a secure environment that uses digital signatures and certificates of authority
- Integrates the address book with directory searching
- Offers a customized user interface
- Supports other protocols for mail access, such as POP3 and IMAP4, and directory usage, such as Lightweight Directory Access Protocol (LDAP)

Disadvantages

Potential disadvantages of deploying Outlook 2000 include:

- Runs only on Windows 95, Windows 98, Windows NT or Windows 2000 operating systems.
- Bandwidth intensive from a protocol perspective (requires custom ports, open firewalls).
- Requires installation and drive space on the local workstation.
- Requires extensive training time due to the multiple functions supported.
- Must be distributed and managed just as any distributed and locally installed software. This means you must purchase either multiple copies of the software or end-user licenses, you must manage licenses when users leave or are hired, and you must either configure each client or provide installation instructions.

IMAP4 and POP3

IMAP4 and POP3 are Internet standard protocols for retrieving mail from a mail server. POP3 provides less functionality and requires less administrative configuration than IMAP4. It is designed primarily for offline use. POP3 allows users to connect to a server and download all of their e-mail from their inbox to the client. POP3 presumes that all mail is retrieved from the mail server, sent to a local mail store, and then kept at the local mail store. The IMAP4 protocol, which provides more functionality and requires slightly more configuration than POP3, enables clients to access and manipulate messages stored within mailboxes and public folders on the server, instead of managing mail on the client computer. IMAP4 assumes that all the messages are kept on the mail server and allows the user to maintain a local view of the folders and messages stored on the server. Both rely on SMTP to send messages and LDAP for directory access.

POP3 Advantages

Advantages of deploying a POP3 client include:

- Offers basic functionality, requiring minimal configuration, and it is a fast protocol
- Provides an offline solution for reading and sending messages
- Most firewalls that are enabled for mail allow the POP3 protocol to pass through by default
- Is a widely-used standard on the Internet
- Minimizes user downtime during active use, except when synchronizing mail received or sent

IMAP4 Advantages

Advantages of deploying an IMAP4 client include:

- Allows the user to read messages based on the folder infrastructure on the mail server
- Allows management of mail-based folders and messages
- Provides offline capability by either maintaining a copy of the mail server folder structure or by allowing the user to copy messages to a local message store on the workstation, depending on the messaging client—Outlook Express or Netscape Messenger
- Allows access to the public folders
- Further minimizes user downtime during active usage. except when synchronizing received or sent mail
- Allows message headers only to be downloaded initially to reduce synchronization time

POP3 Disadvantages

Some potential disadvantages of deploying a POP3 client include:

- No capability to manage folders on the mail server except for message deletion
- No access to any folder on the mail server except for the Inbox
- Sent messages are not retained on the mail server, but on the local message store
- Messages are typically not stored centrally (although they can be configured to do so); this is typically not acceptable in a transaction-based environment or for legal or accounting usage
- Except for rich clients, such as Outlook 2000, clients typically do not support collaboration software or integration of portal functionality
- Upgrades and service packs must be distributed to every workstation

IMAP4 Disadvantages

Some potential disadvantages of deploying an IMAP4 client include:

- Does not allow for an intuitive interface for non-message based folders such as calendar information, contacts, and so forth
- Depending on the client, sent messages are not stored on the mail server
- Except for rich clients, such as Outlook 2000, clients typically do not support collaboration software or integration of portal functionality
- Upgrades and service packs must be distributed to every workstation

Outlook Web Access

When users read mail from their browsers. Outlook Web Access provides HTML-rendered Web pages that allow access to authorized mailbox and public folder stores. Outlook Web Access presents an Outlook-type appearance to browser-based users. Functionality varies slightly depending on the choice of browser; it is optimized to work with Microsoft Internet Explorer 5.0. The version of Outlook Web Access included with Exchange 2000 works on an Exchange 2000 server with Microsoft Internet Inforination Services 5.0 (IIS 5.0). Outlook Web Access is also designed to take advantage of some of the server-based scalability options, such as front-end protocol servers that provide a unified namespace for all Exchange servers in a front-endback-end configuration.

Outlook Web Access Advantages

Advantages of deploying Outlook Web Access include:

- Provides an Outlook-type experience for browser-based clients
- Allows Exchange 2000 and Outlook 2000 to be used on a variety of hardware and operating systems
- Requires minimal space for configuration and still has an intuitive interface for access to specialized folders such as folders that store appointments on calendars, and folders that store contact information
- Provides URL access to the private message store
- Client hardware requirements are minimal because all that is required is a browser
- Outlook Web Access-based upgrades and service packs are centralized

Outlook Web Access Disadvantages

Some potential disadvantages of deploying Outlook Web Access include:

- Does not support spell check, although there are browsers that support spell check independently of Outlook Web Access
- No real offline capabilities, although Internet Explorer 5.0 does have some support for offline Web pages
- Does not support the same full-feature set as Outlook 2000
- User experience might vary depending on the browser

Terminal Services

Terminal Services provides terminal access to a Windows NT server for any operating system that can run a Windows Terminal Services Client, including Windows CE. Thus a user running Terminal Services Client obtains a Windows NT session, allowing the user to simulate Windows NT on their local workstation. As such, it is possible to run Microsoft Outlook or any of the other client options. This solution relies on a Windows NT Terminal Services infrastructure.

Advantages

Advantages of deploying Terminal Services include:

- Provides choice of client solutions
- Can cross multiple platforms that run Terminal Services clients
- Provides all the security of Windows NT
- Client upgrades and services packs are centralized to the Terminal Services servers

Disadvantages

Some potential disadvantages of deploying Terminal Services include:

- For non-Windows-based users, training costs are increased
- The number of users supported per Terminal Services is limited
- This is a network-bandwidth-intensive solution
- Requires a separate Windows NT infrastructure to be maintained to support mail
- It is costly

Table 21.1 compares the available functionality of each client. You can use the table to quickly determine which client option provides the functionality that your organization requires.

Table 21.1 A comparison of functionality among Exchange-supported clients

Function	Outlook 2000	IMAP4/POP3	Outlook Web Access	Terminal Services
Calendar scheduling	Yes	No (Outlook supports publishing free/busy information to a URL)	Yes	Yes
Free/busy	Yes	No (Outlook supports publishing free/busy information to a URL)	Yes	Yes
Contact management	Yes	Local only—client dependent	Yes	Yes
PIM	Yes	Client dependent	Partial (Tasks but not Notes or Journal)	Yes
Integration with Office 2000	Yes	Client dependent	Limited, based on choice of browser	Yes
Offline capability	Yes	Client or protocol dependent	No (very minimal with IE5 offline capability)	No
Rich-text support	Yes	Client dependent	Yes	Yes
Spell check	Yes	Client dependent	No	Yes
Inbox rules	Yes	No	Through Office 2000 only	Yes
Notes	Yes	POP3: No IMAP4: read only	No	Yes
Journals	Yes	POP3: No IMAP4: read only	No	Yes

Table 21.1 A comparison of functionality among Exchange-supported clients *(continued)*

Function	Outlook 2000	IMAP4/POP3	Outlook Web Access	Terminal Services
Public folder access	Yes	POP3: No IMAP4: read only	Yes	Yes
Server folder management	Yes	POP3: No IMAP4: Yes	Yes	Yes
S/MIME encryption/signing	Yes	Client dependent	Client dependent	Yes
lntegrated with Microsoft Office	Yes	No	No	Yes
Supports non-Microsoft operating systemc	No	Client dependent	Yes	Client dependent
Integrated address book	Yes	Partially through LDAP	Yes	Yes
Integrated Directory Search	Yes	Partially through LDAP	Yes	Yes
Supports URL Access	Yes	No	Yes	Outlook 2000 only
Supports Collaboration	Yes	No	Yes	Yes
Minimal Workstation Requirements	No	Client dependent	Yes	Client dependent

CHAPTER 22

Corporate Backbone Scenario

Sasha Frljanic, Consultant, Microsoft

This chapter describes how to construct and deploy a Microsoft Exchange 2000 Server message backbone in a corporate scenario. This backbone is designed to carry not only Exchange messages. but messages from other e-mail systems. Using connectors and directory synchronization, Exchange can function as the mediator among various messaging systems. This type of backbone is particularly valuable in environments that are very dynamic, with systems being integrated and separated frequently.

In This Chapter

Financial Bank, Inc.

Using Exchange 2000 as a Backbone

Auditing Your System

Building the Directory Topology

Building the Messaging Backbone Topology

Summary

Financial Bank, Inc.

For the purpose of this chapter, a fictional multi-national corporation will be introduced. This corporation, called Financial Bank, Inc., has a variety of e-mail systems that span many different locations around the world. Financial Bank, Inc. has multiple existing Exchange organizations, with locations that are connected with either mesh network (two or more connections between nodes) or hub and spoke topology.

Financial Bank, Inc. has the following e-mail systems:

Microsoft Exchange 5.5: approximately 45,000 seats

- IBM Host–based mail Systems Network Architecture Distribution System (SNADS): 10,000 seats
- Lotus Notes: 5,000 seats
- Lotus cc:Mail: 4,000 seats
- Novell GroupWise: 1,000 seats

The goal is to successfully implement Exchange 2000 Server as an underlying backbone for existing mail and directory infrastructure. The benefit of Exchange 2000 is that it reduces the complexity of the e-mail environment and allows for simplified and centralized management of the e-mail backbone. To ensure the successful deployment of the messaging infrastructure, a design has been chosen that addresses four requirements:

- Seamless national e-mail and directory integration with the Active Directory directory service
- Ease of administration
- Reliability
- Scalability

Using Exchange 2000 as a Backbone

To meet some or all of these requirements, consider the features that come with deployment of Exchange 2000. All of the services discussed in this section, enhance the backbone in one way or another.

Advanced SMTP Command Verbs

Many mail systems, including Exchange 5.5 and Exchange 2000, comply with the Request for Comments (RFC) 821 Simple Mail Transfer Protocol (SMTP) specification. RFC 821 defines the standard SMTP protocol implementation for how to send mail on the Internet. Exchange 2000 supports additional extensions to the SMTP protocol. Two examples are chunking and pipelining.

Chunking

RFC 821 specifies that the SMTP client issue the DATA command to signify the start of the actual message data representation (body or content). This standard also says that the end of the data is signified by a sequence of five characters: carriage return, line feed, period (full stop), carriage return, line feed. This is a slow and inefficient way of sending the data chunk, because the SMTP host has to continually scan for the end of the data (body or content). Although this seems primitive, the majority of SMTP servers on the Internet, including Exchange Server 5.5, support only this method of sending data.

To overcome this performance bottleneck, Exchange 2000 implements the BDAT command from the extended SMTP specification for chunking, as defined in RFC 1830. Essentially, this replaces the standard DATA command with the BDAT command and an argument. The argument specifies the number of bytes in the message chunk that the receiver should expect. All the server needs to do now is count the number of bytes received; when this number equals the value given in the BDAT argument, the server assumes it has received the entire message.

An Exchange 2000 server will advertise that it supports the chunking extension when a client sends an EHLO command. Similarly, when an Exchange 2000 server is acting as an SMTP client, it will use chunking if the server advertises it.

Pipelining

The RFC 821 SMTP standard specifies that an acknowledgment be sent for each SMTP command issued. Most commands normally receive the 250 OK acknowledgement, which means that the last command was successful. This acknowledgment procedure can slow down the overall performance of the message transfer, particularly as SMTP implementations are becoming more reliable. The problem is especially noticeable on high latency networks.

To overcome this performance bottleneck, Exchange 2000 implements pipelining as defined in RFC 2197. This allows multiple commands (such as RCPT TO) to be streamed from the SMTP client to host without waiting for an acknowledgement.

The Web Storage System

The Microsoft Web Storage System is a storage platform that provides a single repository for managing multiple types of unstructured information within one infrastructure. The Web Storage System combines the features and functionality of the file system, the Web, and a collaboration server (such as Exchange Server) through a single, URL–addressable location. You can store, access, and manage information, as well as build and run applications. The Web Storage System is based on the storage technology in earlier versions of Exchange.

The Web Storage System provides important enhancements to the data storage architecture, including the addition of a streaming store. This means that each database (consisting of an .edb file and an .stm file) can contain native content formats that are convened only when necessary.

Storage in Native Internet Format

In Exchange 5.5, after the private information store receives a message, the message is converted into MAPI and stored. When an Internet client, such as Internet Message Access Protocol version 4 (IMAP4) or Post Office Protocol version 3 (POP3), requests the message, the conversion occurs again from MAPI to Multipurpose Internet Mail Extensions (MIME) format. With the Web Storage System, the message is saved in its native format in the streaming database (.stm file); when an IMAP4 request is received, the message does not need to be converted. This reduces the chance of content loss during conversion and improves server performance because conversions occur only when necessary.

Public Folders

Some of the important enhancements in Exchange 2000 are reflected in public folder functionality. The following list highlights some of the changes and improvements to Exchange public folder functionality and technology:

- **Microsoft Management Console (MMC) administration snap-in for public folders** A separate console (Exchange Folders) is available to manage public folders. Therefore, you can manage public folders from a single interface. You no longer need a Microsoft Outlook client to create public folders.
- **Multiple public folder trees** You no longer have to store all public folders in one hierarchy.
- **Exchange Installable File System (IFS)** The IFS allows you to share public folders so a wide variety of clients can be connected.
- **Enhanced security** You can secure items in public folders by using the IFS, and exert more control over the management of a public folder.
- **Accessibility from the Web** Public folders are now easily accessible by using a Web browser and specifying a simple URL for the folder.
- **Full-text indexing** The public folder store has built-in indexing that you control. Outlook clients automatically use this index when performing a Find or Advanced Find.
- **Default referrals** Public folder referrals allow clients to access any folder in the organization. Previously you needed to enable referrals between sites; now referrals are enabled by default between routing groups.
- **Public Folder Inter-organization Replication tool** This replicates information such as Free/Busy System folders, which allows users from different forests to schedule meetings with one another and look up free and busy information.

Be aware that there is no automatic solution for replicating users' calendars between organizations, and therefore users cannot open calendars from outside their Exchange organization's boundaries.

Message Routing

The introduction of intelligent routing in Exchange 2000 complements the improvements provided by SMTP native transport. The Exchange 2000 routing topology is more efficient due to the use of link state information, improved message queuing, and message categorizing.

Link State Algorithm

Although Exchange 2000 uses routes and costs, a new link propagation protocol has been implemented called the link state algorithm. This is based upon Dijkstra's algorithm from 1959 and has been used extensively on the Internet for many years in the form of Open Shortest Path First (OSPF) routers. The link state algorithm is responsible for propagating the state of the messaging system in near real-time to all of the other servers in the organization. This has the following advantages:

- Each Exchange server can make the best routing decision at the source instead of blindly sending a message down a path where a downstream link might be down.
- Each Exchange server can determine whether alternate or redundant links are functioning, which eliminates message bounce between servers.
- It overcomes message–looping problems.

Link state data is propagated between routing groups through SMTP on port 25, and within the routing group using TCP port 3044.

Advanced Queuing Engine

The advanced queuing engine is at the core of Exchange 2000 transport. All messages that are submitted to the Exchange server (including local messages) must pass through the advanced queuing engine.

You will notice two fundamental differences in the way that messages are handled between earlier versions of Exchange Server and Exchange 2000. First, all messages are sent to the transport, even when the sender and recipient are located on the same server. This allows for custom event sinks to operate even in a single server environment. Second, the advanced queuing engine reads the message data directly out of the Web Storage System through a file handle, thereby increasing performance. In earlier versions of Exchange, the message transfer agent (MTA) process would physically copy the data out of the Web Storage System and then transport it. You can see this change for yourself as messages sent from an Outlook client disappear from the Outbox faster than with earlier versions of Exchange.

Message Categorizer

The message categorizer performs lookups and checks limits and restrictions in Active Directory. Additionally, the message categorizer expands all groups. Microsoft Windows 2000 Server includes a basic message categorizer called Cat.dll. Installing Exchange 2000 Server upgrades this component so that it can read Exchange-specific attributes in Active Directory.

Exchange Instant Messaging Service

Exchange Instant Messaging Service is designed for individual users to have one-to-one interactive conversations through their computers. Instant Messaging allows you to send messages and receive an immediate response from the recipient. You can also view presence information, which shows whether other users are online at their computers, out of the office, or not receiving calls. For example, members of a virtual team that are collaborating with one another to produce a single product might need the urgency that Instant Messaging can provide. In an Internet environment, friends can chat with one another on a live discussion thread without the effort of composing and sending e-mail. All Instant Messaging communication takes place over Hypertext Transfer Protocol (HTTP). The message format is Extensible Markup Language (XML). The Instant Messaging server itself runs as part of the Windows 2000 Server Internet Information Services (IIS) process (Inetinfo.exe) and is implemented in Isapi.dll.

Chat Services

You can deploy Exchange 2000 Chat Service in situations in which individuals need to communicate in groups, such as in a discussion forum. Unlike Instant Messaging, which is designed primarily for one-to-one communications, a chat community enables you to join a channel (chat room) and communicate with others in a forum. You can also configure Chat Service to allow users to create dynamic channels, which exist only for the duration of a spontaneous conversation.

Data Conferencing Provider

Data Conferencing Provider enables you to pre-arrange an electronic conference. Participants can share multimedia information and applications, and chat with each other. The Exchange 2000 Data Conferencing Provider also overcomes other problems inherent to Microsoft NetMeeting in the following ways:

- The conference meeting session is scheduled by using Outlook and then started, managed, and closed-down by the Exchange 2000 server running Exchange 2000, so that there is no one single workstation acting as the host.
- If the server running Data Conferencing Provider is located in a perimeter network (also called a demilitarized zone [DMZ]), clients from inside and outside the organization can join the conference without unnecessary risk to company security.

Exchange 2000 Data Conferencing Provider supports all of the T.120 standard conferencing facilities that are provided by NetMeeting. In addition, Exchange 2000 Data Conferencing Provider supports Internet Protocol (IP) Multicast audio and video conferencing. IP Multicast implements a lightweight, session-based communications model that places relatively little burden on conference users. By using IP Multicast, users send only one copy of their information to a group IP address that reaches all recipients. IP Multicast is designed to scale well as the number of participants expands. Unlike NetMeeting, adding one more user does not add a corresponding amount of network traffic. Multicasting also results in a greatly reduced load on the server.

Note Not all routers and IP networks are multicast–capable. This is an important consideration for using Data Conferencing Provider with your existing network infrastructure.

If you have used NetMeeting before, you are probably aware of some of its limitations:

- Most clients do not have a scheduling facility.
- The NetMeeting paradigm works best when only a small nutnber of people want to join a meeting, because all communications are funneled through the host workstation.
- If the host computer drops out of the call, the conference ends.
- The meeting model works in an Exchange organization or over the Internet, but not both simultaneously.
- Nearly all companies connected to the Internet have a firewall in place for security. Although it is possible to allow T.120 sessions through a firewall, it's not feasible. A workstation must host the meeting, and configuring the firewall to allow this traffic directly through to workstations creates too much processing load and potential security risk.

Auditing Your System

In any large Exchange deployment, there is a requirement to assess the actual deployment against the initial deployment plans. Specifically, the audit process is needed to validate the initial recommendations and incorporate them. Auditing may also expose new issues required for final deployment and server configuration. Microsoft Consulting Services conducts audits that focus on the following areas:

- Process of deployment, including a pilot phase
- Existing e-mail infrastructure: message routing and connectivity to other mail systems
- Directory infrastructure
- Performance

Some basic design questions to ask at each phase include the following:

- Should the company upgrade the existing backbone or build a new parallel system?
- Should deployment consist of multiple Exchange organizations or a single organization?
- How do you know if the network can accommodate Exchange 2000?
- Which connectors to other e-mail systems should the company use?
- What are the plans for redundancy?
- How do you know if the system is performing optimally?

The following sections discuss the results of the audit on Financial Bank, Inc., and present some recommendations about improving the backbone. Planning considerations and implementation of these recommendations are covered later in the chapter.

Directory Synchronization

Prior to their Internet mail, Lotus Notes, and Microsoft Exchange Server implementations, Financial Bank, Inc. used the Enterprise Address Book on an IBM host system as the master directory, with all other mail systems receiving updates from the Enterprise Address Book. In addition to this directory implementation, the following directories are also running:

- Internet mail: Internet alias file
- Lotus Notes: Notes Name and Address Books
- Microsoft Exchange: global address list (GAL)

Each of these has its own master directory requiring manual updates. To consolidate all directories, Financial Bank, Inc. requires a mechanism that performs updates and changes automatically.

Message Routing and General Configuration

Many routing mechanisms for e-mail exist between messaging systems. It is recommended that these various routes be removed and centralized in such a way that all e-mail systems send mail using an Exchange 2000 hub responsible for messaging connectivity.

Internet Mail Architecture

Each location has Internet access and the SMTP connector to the Internet as shown in Figure 22.1.

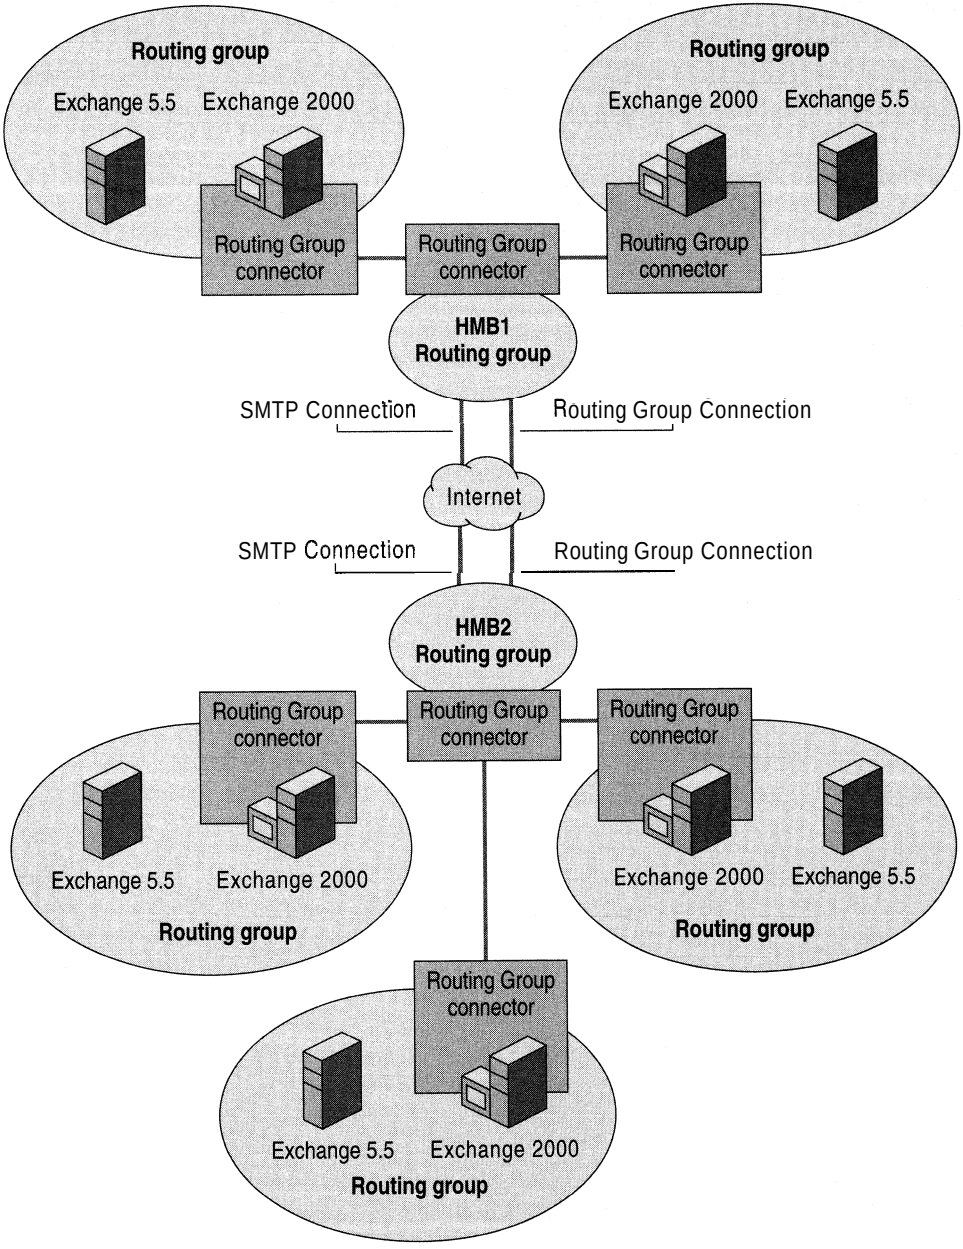


Figure 22.1 Internet mail infrastructure

It is recommended that you consolidate all routes to the Internet and create a redundant, single point of entry from the Internet, which can be monitored around the clock. This solution reduces the overhead associated with maintaining multiple locations. Additionally, it provides increased security and monitoring. It is easier to create a backup plan for a single gateway should something

happen to the primary system. At Financial Bank, Inc., each national implementation will keep their outgoing SMTP routes for a period of 6–12 months. After that period, all incoming and outgoing traffic from the Internet will go through the single point. This will take advantage of using Exchange 2000 to relay mail throughout the company.

Building the Directory Topology

This section discusses some of the concepts involved with building an Exchange 2000 backbone directory topology. It also explains how directory-related deployment decisions are made for the Financial Bank, Inc. deployment.

Planning Active Directory

Financial Bank, Inc. wants to consolidate all directories. Exchange 2000 offers several ways to connect different directories together. The decision has been made to create a new parallel backbone architecture based on Exchange 2000 servers. All locations will gradually migrate to the Exchange 2000 backbone. If Exchange 2000 does not provide a connector for a non-Exchange system, such as Professional Office System (PROFS), Financial Bank, Inc. will use an earlier Exchange 5.5 server to host the connector in the Exchange 2000 organization.

For connectors to other mail systems that are hosted on Exchange 5.5 computers, you must implement Active Directory Connector (ADC).

Active Directory Connector

Active Directory Connector (ADC) is the component responsible for synchronizing Windows 2000 Active Directory with the Exchange Server 5.5 directory service. Keep in mind that regardless of whether Financial Bank, Inc. runs other mail systems, ADC must be deployed for the upgrade and migration phases of Exchange 5.5 to Exchange 2000.

During the design phase, Financial Bank, Inc. must consider a few design issues. First, because they want to manage all mailboxes from Active Directory, they will need to deploy two–way connection agreements that can write to the Exchange 5.5 directory. As the ADC replicates objects from the Exchange 5.5 directory to Active Directory, it must write various attributes into each Exchange object that it touches. Because of this, there is a one-time replication that can take a considerable amount of time, as each Exchange object in the Exchange 5.5 sites are replicated to Active Directory. In very large Exchange systems, this replication can cause directory and network problems.

The other consideration is to determine how many Active Directory domains should be created. When Financial Bank, Inc. deployed Exchange Server 5.5, they had a requirement to move users seamlessly between Exchange servers. To meet this requirement, they deployed very wide Exchange sites that spanned low-bandwidth networks. This resulted in additional management and administration challenges.

For the best performance, Financial Bank, Inc. decided to install ADCs on member servers in the Windows 2000 domain. Multiple ADC servers are deployed across the company and configured with connection agreements so that each replicates a different set of objects.

> **Note** In some cases, it may make sense to install an ADC on a domain controller to reduce network traffic. If you do this, it is important to resize these servers to accommodate the extra processing load.

Each Exchange mailbox is matched to a primary Microsoft Windows NT 4.0 account; such accounts normally reside in a master accounts domain. In the ideal environment, each of these account domains should be upgraded to Windows 2000 Server before the ADC is deployed. This doesn't mean that all servers in all domains need to be upgraded to Windows 2000 Server at once—only the primary domain controller of each domain requires the upgrade. All other backup domain controllers and member servers can reside in the mixed-mode domain.

When considering the deployment strategy for Active Directory, it is important to plan the number of domains, the domain tree hierarchy, and the location of resources. In an ideal situation, the company would deploy a single Active Directory forest to enable a centralized management and security model. However, Financial Bank, Inc. decided to deploy multiple forests for the following reasons:

- Geographical boundaries within the company are to be maintained.
- Legal situations exist in which directories must be completely isolated from one another. Financial Bank, Inc. has some offices that require directory isolation and which are independent of the rest of company.

  Departments do not need to communicate with one another.

However, it is important to remember that only one Active Directory can own a single domain name. In the Financial Bank, Inc. scenario, there has been a communication failure between teams, and an independent department has deployed its own forest. Having multiple forests with one information technology support infrastructure can be a challenge and can affect the deployment of Exchange 2000. In this situation, manual trusts between domains must be created in separate forests. However, because these domains are non-transitive, they decided on a domain model that resembled a Windows NT 4.0 deployment, with multiple manual trusts between each domain. Additionally, the global catalog servers knew about only the objects within their own forest, which had a direct impact on what Outlook users could see in their address book. Thus, a plan was adopted to migrate the independent department to the central forest after the infrastructure is deployed throughout the rest of the company.

Because Exchange Server 5.5 performed mail-based directory replication between sites, users in all domains were able to see all users within the same organization through one global address list (GAL). However, if you consider a similar scenario with Active Directory and Exchange 2000 deployed in separate forests, issues arise because a single Exchange 2000 organization cannot span two forests. The ramifications of this alternative scenario on Exchange 2000 are as follows:

- Two separate Exchange organizations to administer
- No automatic directory replication between the two organizations, which results in two separate GALs and users only being able to see one of the lists
- Cannot use Routing Group connectors between the two organizations; instead, SMTP connectors or X.400 connectors must be used
- No link state data transfer for Exchange 2000 because Routing Group connectors cannot be used

If the two Active Directory domains belong to the same forest, a single Exchange 2000 organization can be formed.

Global Catalog Servers

Financial Bank, Inc. took special care with determining the placement and number of global catalog servers that were required. Global catalog servers contain a copy of every object from every domain in the forest, but only a select set of the attributes from each object. The global catalog enables fast, efficient searches that span the entire forest. The global catalog makes directory structures within a forest transparent to users.

The availability of global catalog servers is crucial to the operation of the directory. When Financial Bank, Inc. was deciding on the number of global catalog servers required to support their Exchange 2000 deployment, they considered the number of clients supported by each global catalog and the number of Exchange 2000 servers in each Windows 2000 site. They also evaluated how many directory searches users performed daily. Finally, they had to plan redundancy for global catalog servers.

When Outlook users want to find a person within the organization, they would normally search the GAL, which represents an aggregation of all messaging recipients in the company. Because Exchange 2000 servers no longer host their own directory service, all data is retrieved from the global catalog servers in Active Directory. Because a global catalog server can support the MAPI address book interfaces, as well as Lightweight Directory Access Protocol (LDAP), Outlook 2000 clients can communicate directly with Active Directory using the same protocol employed by the Exchange Server 5.5 directory service.

For redundancy purposes, Financial Bank, Inc. used the same failover and load distribution rules that were used for individual domain controllers to determine whether additional global catalog servers were necessary in each Windows 2000 site.

If reliability is important, you should plan for global catalog redundancy in each Windows 2000 site. If a Windows 2000 site spans multiple domains, it is also recommended that you configure a global catalog for each domain where Exchange 2000 servers and clients have been deployed.

> **Note** As a general rule, you should plan to deploy one global catalog server for each four Exchange 2000 servers within a Windows 2000 site.

Name Resolution

Financial Bank, Inc.'s next consideration was name resolution. Exchange 2000 is a network-oriented product that relies heavily on its underlying network infrastructure and protocols to carry data. Clients communicate with servers to access collaboration data, and servers communicate with other servers to route messages. All products of this nature require a solid name resolution method.

In Windows NT Server 4.0, the Windows Internet Name Service (WINS) provides network basic input/output system (NetBIOS) name resolution and Domain Name System (DNS) provides Windows Sockets (Winsock)-based name resolution. All earlier versions of Outlook and Exchange clients use, by default, the Winsock layer for communications and DNS as the preferred name resolution technique.

In Windows 2000 Server, other components such as log on and domain validation use DNS as the primary name resolution method. Also, when the bridgehead server that hosts the Routing Group connector receives a message, the bridgehead server tries to resolve the target server's IP address by using the standard SMTP resolution process. That is, the bridgehead server first tries to resolve the target server defined on the Routing Group connector by using DNS mail exchanger (MX) records. If no MX records exist for the target server, which is not unusual, a DNS query for an A (host) resource record for the target server is performed. This means that an A record must exist in DNS for all servers running Exchange. When the Windows 2000 DNS service is used, all servers running Windows 2000, including those servers running Exchange 2000, automatically register A records in DNS. If, for some reason, an A record for the target server is not found, the bridgehead server will try to resolve the IP address by using the NetBIOS name resolution process. In short, Exchange performs MX record resolution for SMTP but usually will resolve the target server by using an A record.

Directory Synchronization Between Forests

A number of different options are available for synchronizing multiple forests. The dirsync control is a standard way for synchronizing directory data from different LDAP providers, and has been submitted as an open standard to the Internet Engineering Task Force (IETF). Microsoft Metadirectory Services (MMS) was used to synchronize directories. Third-party products, such as the Compaq LDAP Directory Synchronization Utility or Siemens' DirX Meta Directory could also be used. For more information about synchronizing disparate directories see "Inter-Organization Replication and Directory Synchronization" in this book.

Building the Messaging Backbone Topology

To meet their requirement in the past, Financial Bank, Inc. had to deploy very wide Exchange sites that spanned low-bandwidth networks. As the Exchange site forced messaging and the directory replication model to use remote procedure calls (RPC) over low-bandwidth links, additional challenges occurred.

The decision has been made to create new parallel backbone architecture based on Exchange 2000 servers. All locations will gradually migrate to the backbone. Internally, the backbone uses advanced routing capabilities in Exchange 2000. Because of the independent department that implemented its own Active Directory, Financial Bank, Inc. is forced to create two Exchange organizations.

The proposal is to migrate all the users from the independent department to the central Exchange organization over time. To synchronize two different Exchange 2000 organizations, Microsoft Metadirectory Services (MMS) will be used. In the beginning of the implementation of the backbone, administrative groups match routing groups one–to–one. When the entire migration to Exchange 2000 is complete, Financial Bank, Inc. will switch to Exchange 2000 native mode, in which the administration of routing groups and administrative groups will be more flexible, and will respond to the different needs of the company.

Routing Groups

An Exchange routing group defines a collection of Exchange 2000 servers that communicate directly with each other. Unlike the constraints on sites in earlier versions of Exchange, routing groups can be created and removed as required, and server membership can be dynamically altered. In other words, the entire routing architecture for an organization can be changed very easily, without reinstallation. As the underlying network infrastructure changes—for example, when new network links are created or upgraded—the Exchange routing network can take maximum advantage of the change. The only prerequisite for this is that the Exchange organization must be running in native mode and all routing groups must be a part of the same administrative group.

The only transport protocol used between Exchange 2000 servers in the same routing group is SMTP. In an Exchange 2000–only environment, RPC connectivity is neither desired nor offered as an option. Where an Exchange 2000 server has joined an existing Exchange site, RPC is used between Exchange 2000 and earlier versions of Exchange.

Link State Algorithm Considerations

All enterprise-messaging systems require a way of passing routing data to one another. In Exchange 2000, when a connector is placed between two routing groups, a cost is associated with the link. A cost represents route preference when multiple routes are available; the lower the cost, the more preferable the route.

Exchange 2000 builds upon the routing architecture in Exchange Server 5.5, where the available routes and costs in the organization are available to any messaging server. Earlier versions of Exchange used a routing calculation server that collected the available routes and costs from the directory and formed a Gateway Address Resolution Table. This table was propagated to all of the other servers in the Exchange site, to provide the routing topology for the organization. However, the routing architecture cannot identify routing problems outside the site. If a network link or bridgehead server is unavailable, there is no mechanism for communicating this information to the rest of the Exchange organization.

Exchange 2000 resolves these issues with changes to directory and messaging architecture. You can create any number of routing groups; however, it is recommended that you limit the number of groups to less than 1,000 for administrative reasons. More routing groups result in larger link state databases, and potentially more status data to replicate; however, this should not seriously affect the performance of the messaging infrastructure. As a guideline, each object in the link state database (routing group, connector, server) requires roughly 32 bytes of memory. Therefore, an Exchange organization with 200 routing groups, 250 connectors, and 500 servers will require just about 32 kilobytes (KB) of memory on each server for the link state database.

Routing Group Deployment Scenarios

Before looking at how Financial Bank, Inc. decided to deploy and connect routing groups, it is helpful to consider how a routing group is defined.

To be included in a routing group, Exchange servers must belong to the same Windows 2000 forest, and must have permanent and direct SMTP connectivity to one another. All servers within a routing group should always be able to contact the routing group master. A routing group might need to be divided when the underlying network between the servers in the routing group experiences frequent reliability problems, or when low-bandwidth connections exist between servers.

Once the routing group boundaries have been defined, the groups need to be connected to one another using a connector. As with earlier versions of Exchange Server, a number of options are available. As Financial Bank, Inc.'s company architects discussed the design, several scenarios came up for connecting routing groups.

Routing Group Connector

The Routing Group connector is the preferred connector for connecting routing groups, because it uses SMTP as the native transport mechanism and obtains its routing and next hop information from the link state database. An advantage of having Routing Group connectors deployed is that the connector can be configured with one or more bridgehead servers on either end of the connector. This provides redundancy should one of the bridgehead servers fail; if this happens, Exchange 2000 can choose another bridgehead server within the routing group to transmit the message.

X.400 and SMTP Connectors

Where there was an intermittent connection or no direct connectivity between routing groups. Financial Bank, Inc. deployed an X.400 or an SMTP connector to enable messaging. Such connectors are appropriate where asynchronous dial-up through a modem is the only form of connectivity between locations. Unlike earlier versions of Exchange, Exchange 2000 does not include a Dynamic RAS Connector. Instead, more efficient routing can occur by using one of the connectors over an on-demand connection supplied by the operating system. This takes advantage of the Routing and Remote Access components included with Windows 2000.

A possible concern is the performance of an SMTP transport over X.400. Although the SMTP transfers can be larger in some cases, this does not necessarily mean that it takes longer to transfer. Tests show that an Internet Mail Service in Exchange Server 5.5 can out-perform the X.400 connector by as much as 300 percent, mainly because both RPC and X.400 transports rely on stringent call setups, handshaking, and acknowledgements. SMTP performance also improves in Exchange 2000 because of the additional performance improvements made in the Exchange 2000 SMTP stack, such as the CHUNKING and PIPELTNING commands.

SMTP connectors are also used to connect servers that perform separate tasks. For example, one server can queue messages while the other server pulls them by using the TRN or ETRN cotnmands. Some remote sites use this type of connector because they do not have high-speed links.

X.400 in Networks Connected with X.25

Financial Bank, Inc. also had sections of their network that were X.25-based. The only way to connect routing groups between X.25–connected networks is to use the X.400 connector. To connect two routing groups with the X.400 connector, you must configure X.400 connectors in both routing groups to create a two-way connection. To provide load balancing, the Exchange designers at Financial Bank, Inc. configured multiple connectors between the two routing groups.

Implementing multiple connectors gives the best results because small text messages will be transferred over fast SMTP, whereas large messages will go over X.400. It's safe to implement parallel connectors between routing groups because they use link state information to calculate the best route. Also, if the network link fails, messages will not bounce back and forth between the connectors if both of these connectors rely on the same physical network link.

Message Routing and Group Expansion

All message routing information, including routing group definitions and bridgehead servers, is held within the configuration naming context of Active Directory. To calculate routes, an Exchange 2000 server contacts a local domain controller and retrieves this information.

If a message is sent to a universal group, the SMTP virtual server, which is configured to perform the expansion, uses LDAP to contact the global catalog server and populate the message header with the group members. If the message is for a domain local or global group, the expansion server uses a local global catalog server, if the membership can be retrieved. For example, if the group is defined within the same domain as the expansion server, then the members can be retrieved. If the members cannot be retrieved because the group is not local to the domain to which Exchange is making global catalog requests, the message is not delivered. Thus, it is very important to set distribution list expansion servers correctly if domain local or global groups are used as distribution lists. To avoid this issue entirely, always use universal groups.

Messaging Coexistence

Although the native transport has changed in Exchange 2000 to SMTP, an instance of the Message Transfer Agent (MTA) still exists in Exchange 2000. There are not too many differences between the MTA in Exchange 5.5 and Exchange 2000.

The single–hop–based nature of an Exchange 5.5 site does not change when an Exchange 2000 server exists in the Exchange 5.5 site. Servers running earlier versions of Exchange are still able to communicate with the MTA by using RPC.

If there are two or more Exchange 2000 servers in an Exchange 5.5 site, the Exchange 2000 servers route messages to one another by using SMTP rather than the MTA process. Thus Financial Bank, Inc. can take advantage of the advanced queuing and routing mechanisms that exist in Exchange 2000. Additionally, as SMTP is asynchronous and can operate in very low-bandwidth network conditions (unlike RPC) it is appropriate for Financial Bank, Inc. to upgrade the hub location dedicated to message switching.

Connectors to Other Mail Systems

Financial Bank, Inc. used earlier Exchange 5.5 systems to connect to their existing IBM Host system. The plan is to migrate existing users of cc:Mail, Lotus Notes, and GroupWise to either Exchange 5.5, or directly to Exchange 2000 servers. After that, Financial Bank, Inc. plans to migrate 10,000 users from the IBM Host system. For now, the Exchange 2000 backbone will be capable of performing all services that are required for the company. However, although Exchange 5.5 supports SNADS and PROFS connectors, Exchange 2000 does not. If you plan to migrate these users to Exchange 2000, you need to use the SNADS and PROFS connector on an Exchange 5.5 server.

Summary

Financial Bank, Inc. has deployed a hub location so that spokes can have a local connection into the message backbone network. In the past, the company had received RPC timeouts and performance problems with this deployment method, due to high variance in network traffic. With Exchange 2000 servers in the Exchange 5.5 sites, this issue is resolved. To minimize risk, the Exchange 2000 servers are installed in the hub location, and the messaging and directory replication connectors will be migrated one at a time.

All locations attach directly to the backbone. All message routing and directory replication between locations use the Exchange 2000 backbone. Each location will implement two connectors to the backbone, with one being designated as the primary connector to the backbone server with the best direct network connection. The server with the primary connector should be dedicated to the routing function, supporting no user mailboxes or public folders. The second connector will be the backup connector and will be connected to a second backbone server. This connector will terminate at the backbone server with next best network connectivity. As more company-wide implementations are added to the global Exchange network, it might be necessary to expand the Exchange 2000 backbone to incorporate other regions.

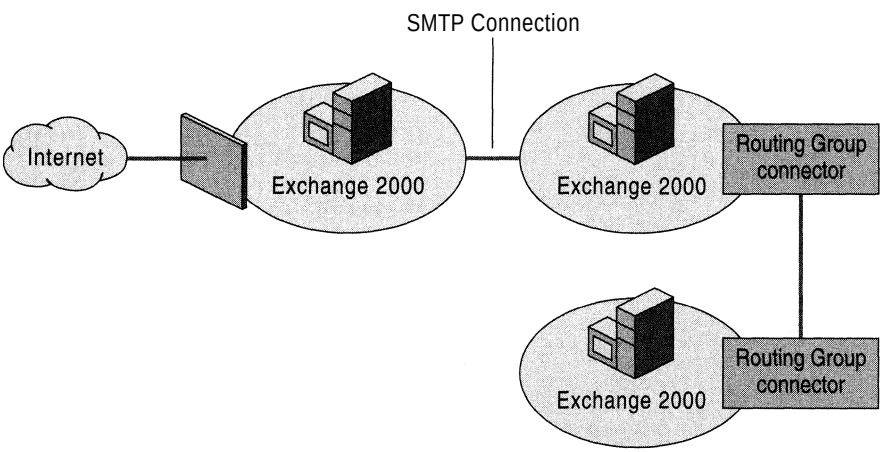


Figure 22.2 Exchange 2000 deployed as backbone

Financial Bank, Inc. consolidated all routes to the Internet and created one point of entry from the Internet. All inbound and outbound traffic from the Internet will go through the one point and use the advantages of Exchange 2000 to relay mail throughout the company.

CHAPTER 23

Hosted Service Environments

David Hitchen, Consultant, Microsoft

Before you host arid provide Exchange services to other companies, you need to understand how hosting can affect your company and the companies that purchase your services. Hosting a service means using Microsoft Windows 2000 Server arid Microsoft Exchange 2000 Server to provide messaging and collaboration services to multiple companies.

Many companies value the features found in Exchange 2000, and Exchange provides you with the flexibility to host services that will fit most of their needs. However, some of the companies to which you provide services may not be in a position to implement Windows 2000 and Exchange 2000. A company may have a relatively small number of users, require further employee training, have financial commitments in other areas, or a need to relieve an overburdened internal Information Technology department.

This chapter provides guidelines and recommended configurations for hosting multiple companies in a single Exchange 2000 environment. This chapter focuses on one scenario that can be expanded upon to provide a more extensive service.

In This Chapter

Hosting Architecture

Service Offering

Customer Division

Administration

Conclusion

Hosting Architecture

This chapter focuses on a scenario of shared hardware hosting that provides the foundation for diversifying into other scenarios; whenever possible, alternative methods are provided.

For shared hardware hosting, there is a single forest that contains all of the user accounts for the hosted companies, and a single Exchange organization that contains all the mailboxes and public folder data. By definition, all objects in a forest exist in a single instance of the Microsoft Active Directory directory service.

In some cases, you may want to host multiple companies in multiple forests. In this case, separate instances of Active Directory are required for each forest (with separate hardware), each with its own Exchange 2000 organization. This is not a technical issue, but one of cost; diversity such as this would require an investment in extended hardware and administration procedures to provide a viable service to customers.

Service Offering

The service you could offer is an out-sourced Exchange service. The focus of the service is a subset of Exchange 2000 Server features. The exact services you offer depend on what you want to provide.

Hosted Exchange services allow a company to take advantage of the Exchange feature set without the overhead of managing and maintaining user accounts, mailboxes, folders, servers, and disaster recovery. The extent to which you manage the service depends on what you offer. For example, you may want the company to administer their own user accounts.

Because Exchange is so versatile, it is impossible to cover all scenarios for hosting Exchange services. However, the following are potentially unique issues for each company for which you host services:

- What client-side operating system is used by the company
- Whether the company uses Microsoft Outlook or Outlook Web Access
- Which public folder capabilities the company wants
- Whether the company needs to support roaming or disconnected users
- Which connection, speed, and trusts should be implemented
- Whether the company maintains a user account database with which you can synchronize
- Whether the company has an existing e-mail service
- The size of the company (This will help you determine whether to provide services on shared hardware or on dedicated hardware.)

Customer Division

After you have defined your role as the host, or service provider, consider how to organize your customer data. The following recommendations are based on industry standards and are a general guideline.

This service will probably be provided over the Internet. You need to consider this in the context of the client services to be provided. For example, MAPI-based services (such as Outlook) will require a Virtual Private Network (VPN), whereas Outlook Web Access, Post Office Protocol (POP), or Internet Message Access Protocol (IMAP) can be provided directly.

Partition the Data

There are two main areas to partition in this type of configuration. The first is Active Directory, which contains the address lists and recipient information. The second area to partition is the Microsoft Web Storage System (public folder stores, mailbox stores, and storage groups), which contains the recipient's public folder data and e-mail.

Windows 2000 and Active Directory

Exchange 2000 services and security rely on Windows 2000 and Active Directory. To host multiple companies, partition Active Directory in a way that ensures the company's data is not accessible to the other companies you host.

When creating an Active Directory design for partitioning, consider the following issues:

- The design must easily restrict the administration of hosted companies to their specific scope. That means that changing settings for one hosted company should not affect other hosted companies.
- Basic Exchange 2000 services, such as e-mail and public folder use, must be available.
- Design for scalability. If the service grows, Active Directory must be expandable. Follow the practices for Active Directory scalability provided with Windows 2000. Determine how many users your system must manage now, and then design a system that can manage more users than currently exist.

Use the following options to assist you in designing partitions for Active Directory:

- **Option 1: Host all companies in a single forest within a single domain** Organizational units partition each company's users. The advantage is that the least hardware is required to handle the users in all companies. This option is easy to manage with a single administration process, and is scalable.

  A disadvantage of this design is the risk of one company's unintended access to another company's data. Be quite careful to ensure that the security risk is minimal.

  An issue to consider is that Windows 2000 adds users and administrators to a default set of security groups, such as the Domain Users group. Disallow this to prevent users from having access to domain resources, such as printers. This will not prevent you from allowing users and administrators to access resources found only in Exchange.

- **Option 2: Hosting companies in a single forest with multiple domains** Each customer stores user data in a separate domain, and the global catalog would provide information from the directory.

  A disadvantage of this is cost. Hosting each company in a separate domain requires at least two domain controllers per company and you must still prevent the problems present with the single domain method. Because security settings and defaults apply to a domain, the impact of a user being added to a Windows 2000 security group, such as Domain Users, is limited to the company in that domain.

- **Option 3: Host each company in a separate forest** Partition each company's users into a separate Active Directory schema. An advantage of this is minimal risk of inadvertently giving one company access to another company's data. Distributing services across hardware offers other advantages. In the event of a failure, you can continue to provide services to other companies while restoring services to the affected company.

  A drawback of this design is cost. To host each company in a separate forest requires at least two domain controllers per customer, separate administration processes, and disaster recovery equipment. You will also have to manage each and every forest as a completely separate entity. This approach is not practical.

  Generally, you provide a unified namespace for each company, such as, www.winery-co-10.com. Within Exchange 2000. front-end server services are not provided for more than one forest; therefore additional front-end servers are required for each company.

Recommended Active Directory Partitioning

Microsoft recommends that you host all companies in a single forest and domain. Implement this by creating separate organizational units for each company. Each organizational unit will contain the company's users, distribution lists, and security groups to restrict or grant access to Active Directory appropriate1y.

As with any Active Directory design, make sure to minimize the risk of unauthorized data access. With this design, the risk increases with the use of Internet access and partitioning security. Some risk is due to the fact that the service uses the Internet as its primary connection. If a company is allowed to log on, it has additional methods for accessing the global address list (GAL). Security groups contained within the organizational units should provide the partitioning security context for the hosted companies. These security groups would consist of users who would assign rights as required. For example, all users will be restricted to the company's organizational units within Active Directory. Some of those users may also have additional rights to create new objects (users).

After the service is configured, a security risk is in the creation of new objects. Processes must be followed, either manual or programmed, to ensure that ample security is applied to each object. Objects without the correct security context may be able to access objects outside the company boundary. For example, users from one company could potentially have access to user information within another company.

Host Services

The service you provide to a set of users is determined by how you partition the data in Exchange 2000. Some services may depend on the client software in use; others may depend on your administration and data recovery model. This section describes a partitioning implementation that expands to more complex models based on your dependencies.

Mail storage in Exchange 2000 is straightforward. Mail is contained within a mailbox database, and there are multiple mailbox databases that can be grouped together in storage groups. However, there are many questions that need to be considered before any storage scheme can be implemented. The following section presents some of those questions and some possible storage solutions.

What Is the Target Company Size?

Before implementing a data storage design, determine whether your service will support companies with few recipients, or large global companies with many recipients. In either case, a consistent methodology enables service reduction or expansion. In this chapter, a company is small, medium, or large based on the number of users or recipients. Small companies have fewer than 100 to 200 users, medium companies have from 200 to 2,000, and large companies have more than 2,000 users. If you are hosting medium–range companies, make the transition from a shared server hosting method to a purely dedicated server method. This suggestion is based on the number of users in the organization and how many mailboxes can exist in a single mailbox store.

Mailbox Storage

A brief look at the mailbox storage options is provided to better design your storage schema Public folder stores are discussed later in this chapter.

Grouping multiple mailbox stores into storage groups provides a natural partitioning design. It is possible to assign each company to their own storage group that contains a mailbox store and a public folder store. This configuration provides a good partitioning service and gives each company assured independence. The databases and transaction logs are independent, as is the administration and disaster recovery.

A major limitation of this design is the small number of recipients allowed per company. For example, imagine that you want to host five companies. Because one Exchange 2000 server can contain only four storage groups, if you assign one company per storage group, you can only host four companies per server. Although only four storage groups are active per server, each with their own transaction log files, you can have five different mailbox and public folder stores per storage group, and as many users as your hardware can handle in a store. This means each company's storage group can contain five mailbox and public folder stores. with a total of 20 mailbox and public folder stores per server. Although this limits the number of companies per server, it is a possible solution for hosting a small number of larger companies, in which users number closer to 1,000.

An alternate, cost–effective service solution for small companies is hosting multiple companies per mailbox store. These companies can share the same mailbox store and storage group. Active Directory and its associated security provides the visible partitioning to the user in the form of GAL and naming schemes. Because the public folder store and mailbox store contents are hidden from the company, except for its own data, little configuration work is needed.

Limitations must meet the company requirements. Determine the total number of users for a single server, or if one company requires large amounts of bandwidth or storage. This has impact on the level of service you can deliver. If a company has its own storage group, backup and recovery is a simple task because service failure affects a single storage group and a single company. However, if there were multiple companies within a mailbox or public folder store, backup and recovery affects all the companies in that mailbox or public folder store. Service disruption may affect an entire storage group and all its stores and companies, and a complete server failure may affect thousands of users across many companies.

Whether you partition each company into its own storage group or partition all companies across a server's 20 mailbox and public folder stores depends on the sets of users. However, both partition designs can be arranged as needed, even on the same server.

Example Service Offerings

This section provides some sample solutions for hosting multiple companies. Before a solution can be provided, the service offering must be defined. For the purpose of our example service, the offering is defined as follows:

- A four-hour recovery in the event of service failure
- Support of MAPI and Hypertext Transfer Protocol (HTTP) clients

In this solution, 150 megabytes (MB) is allotted to each user. Storage size is contiguous storage space: if a company has ten recipients, the total storage space for that company is 1.5 gigabytes (GB), including mailbox stores and public folder stores.

One way to implement this service is to use one server for hosting mailbox data, and another for hosting public folders. Unlike HTTP clients, MAPI clients view one top-level folder hierarchy, therefore additional public folder hierarchies are not available to MAPI clients.

Determining the maximum number of recipients per server depends on many factors, including hardware and recipient usage patterns. The most current hardware enables fast data backup and recovery, failover of services, and the use of Storage Area Network–based data centers. The discussion of hardware is outside the scope of this chapter, however, some assumptions are:

- Servers have failover capability.
- Recovery performance of 50 MB per second is possible.

The use of failover technology reduces the possibility of a total service failure. Because all servers have failover capability to recover the entire server, the focus of data recovery in this example is on database recovery. Table 23.1 contains the potential recovery times per database. In instances when all recipients used their maximum storage allowance, primarily in the mailbox store, the time to recover a database is calculated.

Table 23.1 Database recovery times for a sample service

Users	Mailbox Usage	Public Folder Usage	Mailbox	Time to recover
500	150 MB	0 MB	Approxiinately 73 GB	Approximately 1 hour, 30 minutes
1,000	150 MB	0 MB	Approximately 147 GB	Approximately 3 hours, 0 minutes
1,250	150 MB	0 MB	Approximately 183 GB	Approximately 3 hours, 40 minutes
1,500	150 MB	OMB	Approximately 220 GB	Approximately 4 hours, 30 minutes

In this example, the recovery time is four hours. To achieve this, target a three–hour maximum recovery time, because other time intensive operations are included. It takes time to identify a need for recovery, to retrieve the tape, and so forth. These time scales are worst–case scenarios. Recipients would probably not store all their data in the mailbox database, and a company would probably not use 100 percent of their store allowance.

The data in Table 23.1 illustrates that it is possible to provide 1,000 recipients per database with a four-hour database recovery. Exchange 2000 supports four storage groups, each with five databases of 1,000 users. This allows a hosting maximum of 20,000 recipients. In practice, the total number of recipients per server depends on usage. The actual number of recipients per server can be limited to a much smaller number, perhaps, 6,000. The actual number of recipients depends on the user profiles. Therefore, if your hosting profile is similar to MSN Hotmail or involves heavier usage than that, use a much smaller disk quota, and you may achieve 50,000 recipients per server.

Although partitioning user data is a key resolution factor before the service goes forward, other issues may affect your service, and are discussed in the following sections. They are: naming conventions, GAL, public folders, shared services, and desktop services.

Naming Conventions

When you host multiple companies, design or choose a naming convention that prevents object duplication. Ensure that organization and server names are unique and recognizable to the hosted companies. Check that all users in each company have logon names unique to that company, and to all companies in the forest.

The Exchange 2000 organization name should be unambiguous, because many companies will be using the same service. Names such as Exchange or ASPMail are not recommended. Similarly, label servers in a similar manner. Use names such as msg-01 or msg-02, ending them in a simple sequential numbering system.

Recipient names become more complex in an environment that hosts multiple companies. In Active Directory there are two parts to a user object that must be unique throughout the entire forest: the Active Directory universal principal name, and the earlier Windows name. The universal principal name is used by Windows 2000 to allow a simple universal log on. An example logon user name might be the same as the user's SMTP address. For example, Suzan Fines of Winery Inc. logs on as suzanf@winery.com. The universal principal name logon is supported by Windows 2000. The earlier Windows name is used to log on to versions of Windows that do not support universal principal name. This includes Microsoft Windows NT version 4.0 and Microsoft Windows 98. Like universal principal names, earlier Windows names must be unique within the forest. However, earlier Windows names are limited to 20 characters, which can cause difficulty providing unique names.

Recipients require a unique SMTP address that is unique across the forest and the Internet. This guideline provides a naming convention to use internally. Universal principal name gives you a unique namespace per organizational unit. For example, you can support suzanf@winery.com and suzanf@airlines-international.com without problems.

As a service provider, there are multiple user capabilities to plan for. Universal principal name names and Windows 2000 allow you to use 256 characters, and support domain suffixes to provide an ideal naming service for a hosting environment. In the short term, it is expected that the majority of logon traffic will be from earlier versions of systems where the naming capabilities are not as flexible. Earlier Windows names allow a 20-character maximum, with no domain suffix.

Table 23.2 shows the SMTP name to create for each user that will send e-mail across the Internet. The SMTP name creates a unique user log on for your domain. You can create a universal principal name or name that is appropriate for an earlier version of Windows, depending on the type of operating system to which the user logs on. Make a universal principal name for users logging on to Windows 2000, and an appropriate name for users logging onto earlier versions of Windows.

In the following table, the universal principal name is the same as the SMTP address, excluding the .com or any other first–level designation.

Table 23.2 Examples of user aliases or logon names

Name Used For	User Alias or Log on Name
SMTP	Suzan.Fines\winery-co
Universal principal name	Suzan.Fines\winery-co
Earlier versions of Windows	SuzanF\wnry

It is expected that each company has a domain suffix—their Internet namespace. Using a universal principal name fulfills many recommendations, including: focusing the logon on the hosted company, providing uniqueness within Active Directory, and providing a virtual organization for Domain Name System (DNS) mail exchanger records. A final recommendation is that the name format creates uniqueness within the company's namespace. Use a naming convention that is consistent within an organizational unit and that prevents object duplication.

Global Address List

Active Directory contains a database of all users. The practical implementation of this database, when accessed from an Outlook client, is referred to as the Exchange GAL. When a user searches this list, you want to limit the results to recipients in the same company.

This partitioning of the global address list is supported by access control lists associated with global group memberships. For each organizational unit hosting a company, a global group can be created. The global group membership is the users. Use this group with other permission changes, providing a view of Active Directory specific to the organizational unit's company.

When partitioning the global address list, remove the existing default rights to see all users and objects within Active Directory. Remove from each organizational unit the permissions to the Authenticated Users rights and the Everyone group, if it exists. The host company needs to see all objects within Active Directory.

Add rights for the group account for each hosted company, thus allowing users to see other users in their own organizational unit. Go through each organizational unit and add the group for that hosted company. For example, for the organizational unit representing Winery Inc., add the WineryInc group account and assign Read rights.

There are some permissions that apply only to specific clients. Your service determines which clients get support. For example, Outlook Web Access does not incorporate user permission sets when doing searches. The permissions set on organizational units does not prevent search results from including other company's recipients. If you are using Outlook Web Access, set the attribute *msExchQueryBaseDN* on each address list or user object to limit search results.

As with all changes to the permissions or security, verification and testing must ensure the expected result. In this example, Outlook 2000 or Outlook Web Access gets tested to guarantee that each company views only their company-specific data.

Public Folders

Public folders can have similar restrictions and be partitioned in a similar way to mailbox databases. Exchange supports up to four storage groups that can contain up to five databases that are either mailbox or public folder stores. Just as you can dedicate a server for mailbox stores, you can also dedicate a server for public folder stores. The following discussion pertains to dedicated public folder servers.

The client you choose to provide largely determines the partitioning of public folders. The following section compares the features of Outlook 2000 and Outlook Web Access.

Outlook 2000 is the premier client for Exchange 2000 services. Outlook uses the most Exchange functionality and provides the best performance for the least bandwidth and cost. Outlook 2000 is also a MAPI-based client; it can only view the default public folder tree. If you create new top-level hierarchies to store public folder data in separate databases, MAPI clients do not use the folders in these new hierarchies. Other clients, such as Web browsers, can use the folders in these new hierarchies. The new public folder hierarchies are designed for application access, so if any of your companies use Microsoft Office, for example, they can use a feature called Web *folders* with which you can access these hierarchies.

Outlook Web Access accesses the new public folder hierarchies and provides a subset of Outlook functionality. Microsoft Internet Explorer 5.0 provides the Outlook style interface. The functionality differences between Outlook and Outlook Web Access are detailed in Table 23.3. For a complete feature list, see http://www.microsoft.com/exchange.

Table 23.3 Outlook 2000 and Outlook Web Access features

Outlook 2000 Features	Outlook Web Access Features
Messaging	Messaging
Newsgroups	N/A
Unified inbox	Unified inbox
Creating embedded audio	Creating embedded audio
Creating embedded video	Creating embedded video
Creating voice mail	Creating voice mail
Conferencing and scheduling	N/A
Calendar views	Calendar views
Streaming media	N/A
Default public folders	Default public folders
N/A	New public folder hierarchies

There are benefits and drawbacks to using Outlook Web Access and Outlook 2000. Outlook Web Access offers public folder database partitioning, but reduced functionality. Outlook 2000 has greater functionality and bandwidth, but no access to non-MAPI public folder hierarchies, and decreased server–partitioning configuration. However, Outlook 2000 is more feature–rich than Outlook Web Access. Consider the goals for your specific service to determine which clients to support. When partitioning public folder data, the server's default public folder database is shared by each company and accessed by Outlook 2000 and Outlook Web Access.

Configuring the Default Hierarchy

Because Outlook uses only the folders in the default public folder hierarchy, make a top-level folder for each company under the default public folder hierarchy. Set permissions on the top-level folders so that each company can see their own top-level folder and any required Internet newsgroup folders.

Setting public folder permissions is very similar to setting mailbox permissions. First, remove the Folder Visible permissions from any object, except the administrator. Then add the hosted company's group account to the permissions list and ensure that Folder Visible is set for that account. If some companies require a more detailed permissions set, use NTFS file system permissions to create these permissions.

You can control access at the file system level by modifying the security permissions for the top-level folders. If you assign permissions to the hosted company's group account, all users who belong to the group account can have access.

There are additional services and options within Exchange 2000 and Windows 2000 that require configuration. However, these are discussed in later sections.

Provide Shared Services

There are some services in Exchange 2000 that are configured at an organizational level.

Key Management Services is a security feature that provides keys to encrypt and sign e-mail. However, this option is not recommended for the deployment discussed in this chapter. If you have followed the recommendations so far, hosting multiple companies on shared hardware presents difficulties with Key Management Services. These difficulties range from shared public keys among hosted companies, to certification authority sharing. Additionally, if a company leaves your Exchange organization, the encryption key set would be lost for any existing e-mail.

Two components of Windows 2000 to configure are schema and custom attributes. Changes to these components are system wide. Maintaining custom security attributes for shared hardware that hosts multiple companies can present many challenges. Schema changes cannot be reversed. Companies that request a schema change and then stop using your service leave schema changes in Active Directory.

A service you can offer is connections to other e-mail systems. However, these should be avoided. Although you can create temporary connections to aid in the migration of any existing systems, it is less feasible to maintain multiple connections to systems within multiple companies.

Provide Desktop Services

The proposed service design, thus far, provides complete control over the server side of Exchange, backup and recovery, maintenance, and monitoring. This section discusses managing the client side of Exchange. The following guidelines help you manage the service:

- Provide documentation and training for the services you offer.
- Install and configure Outlook. If the company can support Outlook 2000, you can use a pre-configured media to provide a simple and effective way to implement a specific customization of Outlook and any feature control you need.
- Install and configure Outlook Web Access.
- Designate locations for users to save local personal stores. Suggest that this location be on a company file share.
- Archive user data, such as personal stores (.pst files).
- Import any existing personal address books (.pab files).
- If users require offline storage, which is unavailable in Outlook Web Access, determine storage needs and required additional resources. For example, determine the bandwidth required to store offline address book configurations.
- Decide whether to offer connections to other e-mail systems. Although this type of connection is outside the scope of this chapter, there are some general guidelines to consider. If you provide connections to other e-mail systems, you may affect Exchange on a system-wide basis.

Administration

The administration process is the key to the stability of your service. There are issues such as how to administer passwords and new accounts, changes to existing services, and mailbox sizing. The service example in this chapter provides you with full control over Active Directory and server side services, such as modifications to accounts, public folders, and groups. You can also give some administration control to your customers.

If you expand on this model after initial implementation, follow a planned process for allowing hosted companies to make modifications. Create a central administration mailbox or public folder for companies to process all user requests. For more information on hosting with Exchange, see the Exchange Web site at http://www.microsoft.com/exchange.

Configuration Checklist

Before you monitor the company's use of the service, verify that all of the essential components are configured correctly. Use Table 23.4 as a checklist and refer to the other sections in this book for a more detailed explanation of configuring a given component.

Table 23.4 lists the Windows 2000 Server and Exchange 2000 Server components to configure. The steps to complete are in the Configuration Checklist column. Any additional information appears in the Configuration Notes column.

Table 23.4 Configuration checklist

Area	Configuration Checklist	Configuration Notes
Services	Internet connectivity DNS registration	DNS registration of your company or the hosted company depends on the naming convention you choose.
Active Directory	Organizational units	Active Directory partitioning takes place at an organizational unit level. This is the recommended model for separating each company's data and administration.
Web Storage System	Storage groups Mailbox stores Public folder stores	Use the default storage group on the first server to manage mailboxes and public folders for your own company. You can use all other storage groups for customer company data.
Address lists	Security settings	When a user searches these lists, limit the results to users in the same company. To create this separation, make a global group and then change the user's Active Directory permissions.

Conclusion

You can use existing features in Windows 2000 Server and Exchange 2000 Server to host multiple organizations on a shared hardware platform. Exchange services can be offered in a scalable, reliable, fast, and feature-rich environment to any company with appropriate access. You can expand or customize the basic configuration presented in this chapter to provide the type of hosting scenario that suits your needs and resources.

CHAPTER 24

Security Sensitive Environments

Jung-Uh Yang, Senior Consultant, Microsoft
Matthias Leibmann, Program Manager, Microsoft

Information is one of the most important assets of any company. The most popular application for sharing information in the corporate environment today is electronic mail. Although some e-mail that flows through the corporation might contain relatively harmless information, some e-mail contains highly sensitive information to which no one, other than the intended recipients, should have access.

This chapter presents strategies for exchanging secure e-mail within and between companies. These strategies include using Secure Multipurpose Internet Mail Extensions (S/MIME) protocol between clients, public key infrastructure (PKI) trusts between companies, Windows 2000 Certificate Services, and Microsoft Exchange 2000 Key Management Service. Public key infrastructure can be extended to support various clients and different platforms.

Information about certificates, keys, and encryption is included first, followed by typical scenarios and specific Exchange 2000 Server deployment examples.

This chapter does not cover the following topics:

- Pretty Good Privacy (PGP). This chapter focuses on S/MIME; however, some of the S/MIME considerations might apply and might be useful in a PGP environment.
- Microsoft Exchange Server 5.5 Key Management Service migration to a Windows 2000 PKI infrastructure with Exchange 2000 Key Management.
- Scalability and organizational issues.

In This Chapter

Digital Encryption and Signatures

S/MIME Design Scenarios

Deployment Scenarios

Digital Encryption and Signatures

You can secure e-mail traffic in two ways: through session-based security or through message-based security. Session-based security technologies exploit different messaging layer technologies, such as Secure Sockets Layer (SSL) encryption. In session-based security, you establish a virtual secured channel between the client and the server. Before information flows between the two systems, authentication must be verified. The information that flows in the channel between the two systems is encrypted.

Message-based security secures the e-mail message itself, not just the channel through which the e-mail flows. In cases where highly secure e-mail is essential, message-based security is preferred to session-based security, although both can be used together.

S/MIME is a popular message-based security technology that ensures privacy, integrity, and reliability of sensitive e-mail messages.

The fundamental building blocks for establishing a secure e-mail environment include Windows 2000, Exchange 2000, and e-mail clients. Exchange 2000 provides many ways to enable message-based security, such as SSL client connectivity, encrypted remote procedure calls (RPC), and Key Management Service. Both session-based security (SSL and RPC) and message-based security (S/MIME) are supported through Microsoft e-mail clients. Establishing S/MIME and SSL services requires a PKI.

PKI is generally used to describe the legal topics, policies, standards, and software that regulate or manipulate certificates and public and private keys. In practice, it is a system of digital certificates, certification authorities (CA), and other registration authorities that verify and authenticate the validity of each party involved in an electronic transaction. Standards for PKI are still evolving, even though they are being widely implemented as a necessary element of electronic commerce.

Windows 2000 Certificate Services establishes the PKI. Exchange 2000 Key Management Service is used for the deployment and management of Outlook as an S/MIME client. The e-mail client encrypts and digitally signs e-mail messages.

> **Note** Before reading this chapter, you should be familiar with basic public key technologies and with the principles of certificates and certification authorities. For more information about Certificate Server 2.0 technology, see the section on Windows 2000 security on the Microsoft Web site at http://www.microsoft.com.

This section covers general concepts of e-mail security. The following topics are covered:

- E-mail security
- Certificates and certification authorities (CA)
- Key Management Service
- Key Management Service process flow

E-mail Security

Digital signatures and digital signing are based on encryption technologies. Exchange uses different key pairs for each task in its implementation of these technologies. This section explains how Exchange uses the keys. and how they are generated and stored in the Active Directory directory service.

Message Body Encryption

Message body encryption creates a completely unreadable message body and digital signature. Figure 24.1 shows how the encryption process generates a one-time symmetric key (also called a session key) that encrypts the message body. The symmetric key is then encrypted with each recipient's public key, so that only the recipient can decrypt the symmetric key. The message is then sent. On the recipient end, the private key is used to decrypt the symmetric key, which is then used to decrypt the message. This process is transparent to the user. It is performed with no interaction between clients. Aside from a directory query to obtain the recipient's public key, there is no additional interaction with network services. Symmetric encryption is much faster than asymmetric encryption. This is because symmetric encryption encrypts the data (message body) in bulk.

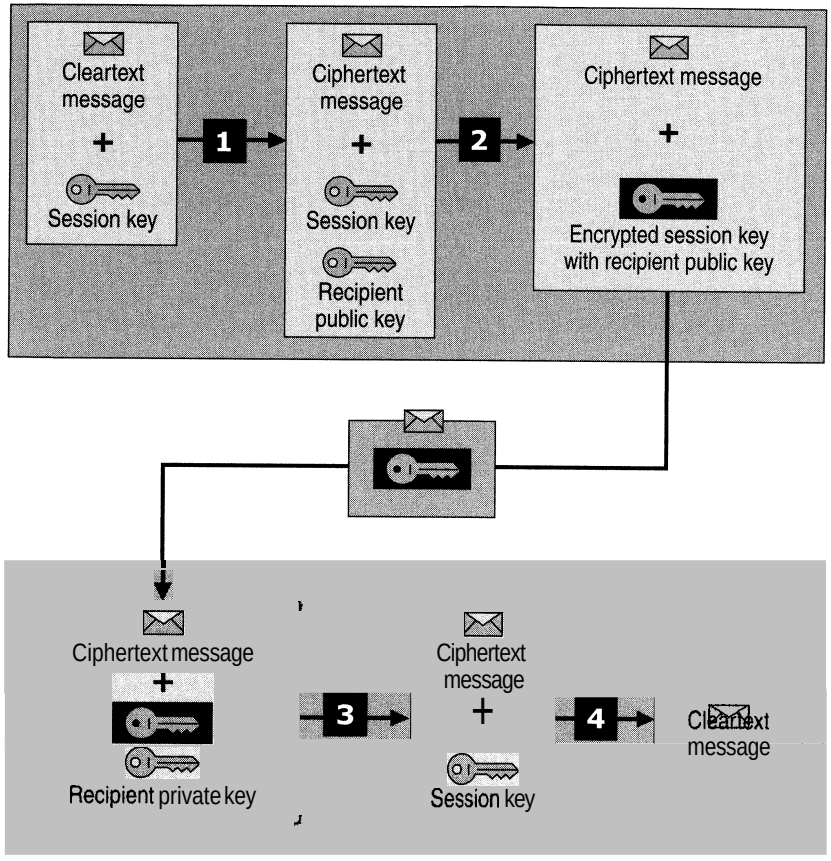


Figure 24.1 Message encryption and decryption process

Figure 24.1 illustrates the message encryption and decryption process. The four main steps detailed in the illustration follow.

1. Message is encrypted with session key.
2. Session key is encrypted with recipient's public key.
3. After encrypted message is received; recipient decrypts session key with the recipient's private key.
4. Message is decrypted with session key.

Digital Signatures

When you add a digital signature to a message, a hash value of the message contents is computed. User keys are not used to compute the hash value, and the hash does not identify anyone. The hash is only a small, unique digital fingerprint of the message. This hash value is then encrypted with the sender's private key, and can be decrypted with the public key found in the sender's certificate. The recipient decrypts the original hash value with the sender's public key, which might be sent with the signed message or can be found in the sender's certificate. You can obtain the sender's certificate from a directory you trust, such as Active Directory.

Your client verifies signatures. For encrypted messages it decrypts the message first. The client computes a new hash value based on the text received and compared to the original hash value. If they match, you can trust the content's integrity and the sender's identity.

Certificates and Certification Authorities

A certificate is a digitally-signed statement containing a public key and the name of the owner (also called the subject of a certificate). The certificate may contain multiple types of names by which the subject is known, such as their directory name, e-mail name, Domain Name System (DNS) name, and so on. By signing the certificate, the CA attests that the private key associated with the public key in the certificate is in the possession of the subject named in the certificate. Certificates may also contain the following additional information:

- The dates between which the certificate is valid.
- A serial number that is guaranteed to be unique to the CA. Serial numbers are guaranteed to be unique only among certificates issued by one CA. If multiple CAs exist, CA1 and CA2 can each issue a certificate with serial number 001, for example.
- The name of the CA that issued the certificate, and the public key used exclusively to verify the signature on the certificates that the CA has issued. The key used to sign the certificate is the CA's private key and is not contained in the certificate.
- An identifier of the policy that the CA followed to validate the subject.
- The uses of the key pair identified in the certificate.
- The location of the Certificate Revocation List that would indicate whether the certificate has been revoked.

Certificates provide a mechanism for establishing a relationship between a public key and the entity that owns the corresponding private key. The most common form of certificates in use today is based on the International Telecommunication Union (ITU-T) X.509 version 3 standard, which is a fundamental technology used in the Windows 2000 PKI. This is, however, not the only form of certificate. PGP secure e-mail, for example, relies on a form of certificates that is unique to PGP.

Certification Authority

A CA is a trusted entity that issues certificates to other entities. A CA acts as a guarantor of the binding between the public key and the subject identity contained in the certificate. Different CAs might choose to verify the binding between the public key and subject through different methods, so it is important to understand the CA's policies and procedures before choosing to trust that authority to guarantee public keys.

The CA is responsible for issuing certificates based on a set of established criteria. The criteria that a CA uses when processing a request is called a CA *policy,* which is different from the general term *policy* that is commonly associated with Windows 2000 domain accounts and application deployment services like Zero Administration for Windows. The CA's policy is typically published in a document known as a certification practice statement.

Three types of certification authorities exist:

- **Self-signed certification authorities** are CAs where the public key in the certificate and the key used to verify the certificate are the same, and the issuer and subject of the certificate are the same. A root CA is a self-signed CA.
- **Subordinate certification authority** is a CA where the public key in the certificate and the key used to verify the certificates are different. This occurs when a CA issues a certificate to another CA to build a hierarchy of trusted entities. This should not be confused with *cross-certification*. Cross-certification occurs when a user, corporation, or other entity trusts another entity's issuing CA. For example, Microsoft and NorthWind Traders may cross-certify by sending their root CA certificates to each other. Neither CA certificate is signed with the other's certificate.
- **Root certification authority** is a special class of CA, which is unconditionally trusted by a client. All certificate chains lead to a root CA. A root CA cannot be designated by another entity as a root; it must be self-signed. Administrators can terminate their trust at any point in the chain, but that does not convert the intermediate CA to a root CA, even to that administrator. The intermediate CA is trusted because it appears in a trust. All self-signed CAs must be root CAs. The certificate chain ends at a self-signed CA.

Certificate Enrollment

Certificate enrollment is the procedure that a user follows to request and receive a certificate from a CA. The identity information that the certificate request provides to the CA subsequently becomes part of the issued certificate. The CA processes the request based on a set of criteria that might require some special authentication. If the request is successfully processed, the CA then issues the certificate to the user.

Public key–based technology generally relies on certificates to bind public keys to known entities. The Windows 2000 PKI supports certificate enrollment to both the Microsoft enterprise CA and stand-alone CA, or to third-party CAs such as VeriSign. Enrollment support is implemented in a transport-independent manner and is based on use of industry-standard Public Key Cryptography Standards (PKCS) #10 certificate request messages and PKCS #7 responses containing the resulting certificate or certificate chain. Exchange supports certificates that support RSA Secured keys and signatures, and Digital Signature Algorithm keys and signatures.

Certificate Revocation

Although certificates have a defined lifetime, CAs can reduce the lifetime of a certificate by the process known as certificate revocation. The CA publishes a list of serial numbers of certificates that it considers no longer usable on a Certificate Revocation List (CRL). Like certificates, CRLs have a defined lifetime, but this is typically much shorter than that of a certificate. The lifetime of a CRL is typically 24 hours to 1 week, whereas certificates are usually valid for one to five years. The CA might also provide a reason why the certificate should no longer be used, and a date from which this change of status applies. Some reasons a certificate may be revoked include:

- Key compromise
- CA compromise
- Affiliation changed
- Superseded
- Cessation of operation
- Certificate hold

If the CA has revoked a certificate, it means that the CA is withdrawing its statement about the allowed usage of the key-pair prior to the certificate's normal expiration. After the certificate that has been revoked expires, its entry on the CRL is removed to keep the size of the CRL to a minimum.

During signature verification, applications can check whether a given certificate and key-pair are still completely trustworthy. If not, the application can determine whether the reason or date of the revocation affects the use of the certificate in question. If the certificate is being used to verify a signature and the date on the signature precedes the date of the revocation of the certificate by the CA, the signature may still be considered good.

Key Management Servi

The Exchange Key Management Service is the center of communication between Certificate Services, the administrator, the cryptographic service provider, the Exchange System Attendant service, and the Archive Database. It is the core component for Exchange Advanced Security, which processes and coordinates the flow of all tasks, such as enrollment, recovery, renewal, and so on.

Cryptographic Services

The Microsoft CryptoAPI provides services that enable application developers to add cryptography and certificate management functionality to their Win32 applications.

> **Note** Application development using the Excacsp.dll (Exchange cryptographic service provider) is not supported nor recommended. Applications can use the functions in CryptoAPI without knowing anything about the underlying implementation, in much the same way that an application can use a graphics library without knowing anything about the particular graphics hardware configuration.

The CryptoAPI provides a set of functions that allow applications to encrypt or digitally sign data in a flexible manner, while providing protection for the user's sensitive private key data.

Independent modules known as *cryptographic service providers* perform all cryptographic operations. One cryptographic service provider, the Microsoft RSA Base Provider. is included with Windows 2000 Server. In the United States, you can also download the 128-bit version of the cryptographic service provider at the Microsoft Website or the Windows Update Web site at http://www.microsoft.com. Each cryptographic service provider provides a different implementation of the CryptoAPI. Some provide stronger cryptographic algorithms while others contain hardware components such as smart cards. In addition, some cryptographic service provider may occasionally communicate with users directly, such as when digital signatures are performed using the user's signature private key.

Key Management Service Passwords

The first security measure for protecting Key Management Service is a startup password. For each restart of Key Management Service, the startup password must be typed manually or read from a removable disk. This password should be stored in a secure location.

The second security measure is administrative passwords, which are required to perform tasks in Key Management Service. By default, only one administrative password is required to perform any task. Administrators, however, can set varying levels of security by requiring up to four different administrative passwords. Each task requires approval by two or more administrators to be executed. Adding and removing Key Management Service administrators and configuring the number of passwords required to execute each function are critical management tasks.

Key Management Service Process Flow

Standard Key Management Service operations include enrollment, renewal, recovery, and revocation.

Enrollment

Advanced Security in Exchange 2000 Key Management Service supports two different message-based security models: S/MIME security X.509 version 3 certificates, and the model in earlier security versions of Exchange that is based on X.509 version 1 certificates. The enrollment process differs in that X.509 version 3 uses Windows 2000 Certificate Services for certificate generation. whereas the Key Management Service itself generated the client certificates in the previous version.

S/MIME Enrollment

Enrollment is a three-stage process in which a user becomes enabled for Advanced Security. To enroll a single user, the following process occurs:

- Administrator initiates enrollment
- Key Management Service generates encryption keys.
- Client stores keys and certificates.

Administrator Initiates Enrollment

1. Open **Active Directory Users and Computers,** right-click on a user object, click the **Exchange Features** tab, right-click **Properties,** and then click **E-Mail Security.**
2. Click **Enroll** to complete the enrollment.

   Note To streamline enrollment, all users within a specific mailbox store, server, administrative group, or organization can be enrolled at the same time. In this case, enrollment takes place through Key Management Service in Exchange System Manager.

3. Key Management Service generates a temporary password, called a *token.*
4. The Key Management Service administrator provides the users with the temporary password, either physically or by sending it to the client using the Key Management Service mailbox agent.
5. In a Microsoft e-mail client, the user clicks **Tools,** clicks **Options.** and then clicks the **Security** tab.
6. The user clicks **Get a Digital ID,** and then selects the option for enrolling in Exchange security.
7. The user is prompted for the temporary password and then for a permanent password.
8. The e-mail client generates a signing key-pair and encrypts the public signing key with the permanent password. The key size is based on the cryptography level of the operating system and Internet Explorer, as well as the presence of the Outlook strong cryptographic service provider.
9. The e-mail client sends a secured signing public key to the Exchange System Attendant.
10. The System Attendant forwards the message to Key Management Service.

Key Management Service Generates Encryption Keys

1. Key Management Service calls the cryptographic service provider to generate an encryption key-pair for the user.
2. The cryptographic service provider then archives the private encryption key and returns the data encryption key-pair to Key Management Service.
3. Key Management Service sends an encryption public key certificate request to Certificate Services.
4. Certificate Services returns the certificate to Key Management Service.
5. Key Management Service archives the certificate and requests the signing key certificate from Certificate Services.
6. Certificate Services returns the signing certificate to Key Management Service.
7. Key Management Service archives the signing certificate.
8. Key Management Service sends the signing certificate, encryption certificate, encrypted private encryption key, and the Key Management Service signing certificate to the System Attendant. Key Management Service signs the certificate trust list that contains the CA certificates.

   Clients need only to trust the signing certificate, which signs the certificate trust list. This is why the trust chain includes a certificate trust list and terminates with the CA that signs the Key Management Service certificate.

Client Stores Keys and Certificates

1. System Attendant forwards certificates and key to e-mail client.
2. The e-mail client creates a file that contains the user's certificates and private keys, each encrypted with the user's permanent Key Management Service password. The file type created depends on the client. For example, Outlook 97 creates an .epf file whereas later versions of Outlook put the key records in the registry.
3. The e-mail client publishes the two certificates in Active Directory.
4. As the client creates a certificate history, certificates are added to the same .epf file or registry key record.

Web-Based and MMC S/MIME Enrollment Process

The Windows 2000 CAs can enroll users through a separate optional component known as the Web Client. Web Client is a Web-based application (Active Server Pages [ASP] for Microsoft Internet Information Service [IIS] addressing Certificate Services) that is installed during setup of the Windows 2000 Certificate Services. By default, Web Client is installed on every CA. It can also be installed on a separate Windows 2000–based server with ITS. Web Client allows Windows 2000 CAs to enroll users regardless of the browser and operating systems they are running. Thus it is possible to deploy Web-based enrollment services to support

non-Windows 2000 systems separately from the CA. S/MIME-enabled mail clients can be enrolled using Microsoft Internet Explorer or any other supported Web browser, for example, Netscape Navigator.

You can enroll S/MIME clients with the Certificates console in Microsoft Management Console (MMC). This enrollment method is supported only for clients running Windows 2000 Professional.

In contrast to the Key Management Service enrollment, Web or MMC enrollment lacks key recovery capability. If a user forgets the password, or the client computer experiences a hardware failure, the lost user's key cannot be recovered and the user cannot read encrypted e-mail.

Before enrolling users, appropriate rights must be set for each user to access IIS, Certificate Services, and the certificate templates in Windows 2000.

In Web or MMC enrollment, the client generates both key pairs. This is different than the Exchange 2000 Key Management Service enrollment process, where the encryption key pair is generated within Key Management server and signing key pairs are generated by the client.

Clients can go to http:llhostnamelcertsrv. where *hostname* is the name of the server that begins the Web enrollment process.

For MMC enrollment, you must install the Certificates console, which contains the Certificate Enrollment Wizard.

Enrolling Users with Certificate Enrollment Wizard

1. A user requests a user certificate (single key-pair, used for both signing and encryption) and a user-signing certificate (dual key-pair request, used only for signing).
2. The base or enhanced cryptographic service provider (available with the Windows 2000 High Encryption Pack) generates a single key-pair (used for encryption and signing) or dual key-pairs (one for encryption, the other for signing) on the client.
3. Client sends the public keys in certificate request message format to Certificate Services. In the Web enrollment scenario, IIS passes this request through to Certificate Services.
4. Certificate Services generates the requested certificates.

   Note If Certificate Services is configured as an enterprise CA, the generated certificates will be published to Active Directory. If Certificate Services is configured as a stand-alone CA, certificates are published in the file system.

5. Certificate Services sends the generated certificates to the client. For Web enrollments, Certificate Services sends the certificates to IIS, and IIS passes the certificates to the user.

Single-Key Pairs vs. Dual-Key Pairs

Clients enrolled on the Web through Certificate Services typically generate one key pair on the client and get one certificate. The key pair is used for encryption and signing. Dual-key pairs are also supported by Web enrollment, where two certificate requests are required: one for a user certificate template and another for user signature template. In this case, all key pairs are generated on the client side.

Key Management Service–enrolled clients use different key pairs. One key pair is generated in Key Management Service (encryption-key pair); the other key pair is generated on the client (signing-key pair). The client receives two corresponding S/MIME certificates: a signing certificate, and an encryption certificate.

Dual-key pairs are more secure than single-key pairs because dual-key pairs separate the encryption process from the signing process. However, dual-key pairs are more difficult to manage.

Outlook 2000, Outlook 98, and Outlook Express support both dual-key pairs and single-key pairs. There are no known interoperability issues with these clients. Some Internet mail clients only support single-key pairs.

In the dual-key pair scenario, both private keys for signing and encryption are stored locally in protected storage. For clients enrolled with Key Management Service, private keys are encrypted with an additional password, which is set during the enrollment process. The corresponding private keys are retrieved from protected storage for the signing and the encryption processes.

Certificate Renewal

Renewal is the process by which users seek new certificates to replace their current certificates. Clients must renew their certificates when either their current certificates are about to expire or when a client has determined that a renewal is necessary. To renew a certificate:

1. Client sends a renewal request to the Exchange System Attendant. The System Attendant forwards the request to Key Management Service.
2. Key Management Service creates a new data encryption key pair, gets a certificate for the public key from Certificate Services, and returns the new key pair to the client through the system attendant.
3. Client creates a new signing key pair and sends the public key to Key Management Service through the system attendant. Key Management Service receives the certificate from Certificate Services and returns it to the client through system attendant.
4. Client adds new certificates and keys to .epf file, and Key Management Service adds new security information to the key recovery database.
5. Client publishes the certificates in Active Directory.

Key Recovery

Key recovery is necessary for two reasons:

- Users can lose their keys. This happens when a user forgets his or her password, for example, or the client computer experiences a hardware failure. Recovery allows users to recover their encrypted e-mail. It also helps your organization recover potentially important information. When required by law, administrators can perform a recovery to gain access to users' encrypted e-mail.
- Key recovery is the final step of the export and import process. Users whose keys have been imported from another Key Management server need new keys. When Advanced Security users are exported from their original Key Management server, their certificates are revoked. When they have migrated to a new Key Management aerver, they can still use their old keys to read old encrypted e-mail, but those old keys are now bound to a certificate published to your organization's CRL. However, users need new certificates and corresponding keys to create new encrypted messages.

In key recovery, as in the enrollment process, the user is issued a token. The recovery token is issued in the same way as enrollment tokens, either through an administrator or through e-mail. After typing this recovery token in Outlook, a new signature key pair is created for the user. In addition, Key Management Service returns all of the user's old keys. For imported users, a new encryption key pair is generated as well.

Certificate Revocation

Revocation is the process by which a user or CA's certificates are marked as unfit for use. Certificates can become unfit when a user is no longer trusted, or when an intruder compromises private keys.

> **Note** The Key Management server must have Manage Permissions on every CA server to revoke certificates and keys. Without these rights, Key Management Service is unable to revoke certificates or keys.

To revoke user keys

1. In the Active Directory Users and Computers snap-in, right-click user object and click **Properties.**
2. Click the **Exchange Features** tab.
3. Select **E-mail security** and click **Properties.**
4. Type the passwords for the certificates that you want to revoke and then click **Revoke.**

The passwords are sent to Key Management Service and verified. Key Management Service sends a revocation request to Certificate Services so it can add the user to the CRL. For more information about password policies, see "Key Management Service Passwords" earlier in this chapter.

> **Note** In Exchange **5.5,** Key Management Service retrieves the CRL from Certificate Services and publishes it so that all clients have access to up-to-date revocation information. In Exchange 2000, Certificate Services must publish to a location listed in the CA's CRL distribution point. Exchange also requires that Outlook check the CRLs online. The CRL is no longer published directly to clients because many CAs can be used at the same time.

S/MIME Design Scenarios

This section will cover the following scenarios for implementing S/MIME technology:

- Intra-company scenario
- Inter-company scenario
- Inter-company with trusted third party scenario

After each scenario is introduced, further detail is provided about trust between companies and trust to third parties followed by high-level planning processes for implementing the scenarios. Implementation solutions and installation hints are covered in "Deployment Scenarios," later in this chapter.

> **Important** Before you deploy any of these designs, be sure to identify all services (in addition to e-mail) that will be enforced with PKI services. If you will be implementing other services, such as Web service, message queuing, code signing, or log on, in the same PKI, you must plan your strategy accordingly. Otherwise you run the risk of re-deploying the PKI to accommodate these services.

Intra-Company Scenario

The Intra-company scenario assumes that the entire CA infrastructure exists wholly within a single company. Figure 24.2 shows the elements of the CA infrastructure.

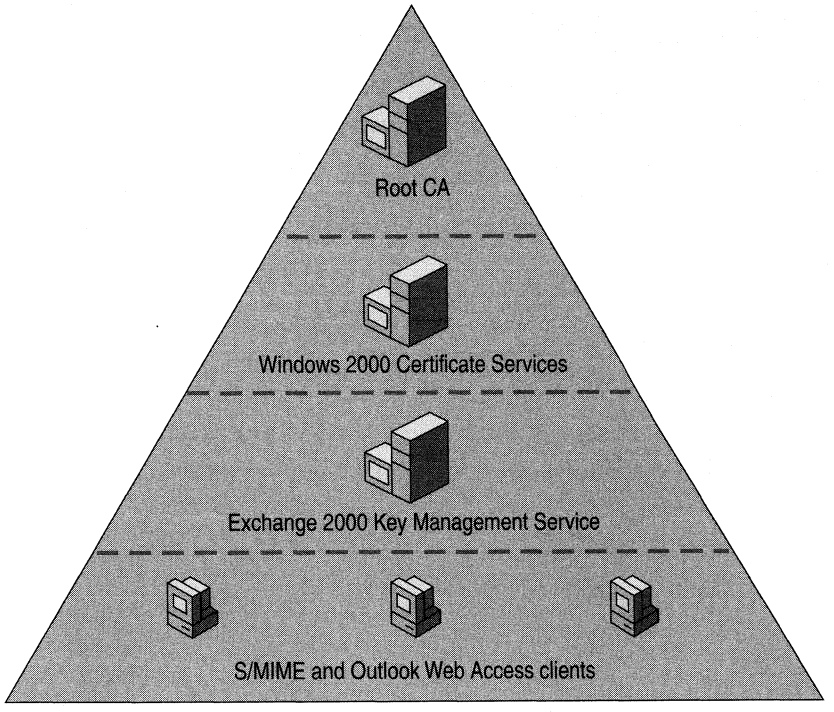


Figure 24.2 Intra-company security infrastructure

Technical issues and features of this scenario include:

- E-mail clients communicate via secured or unsecured e-mail internally
- External e-mail clients and internal e-mail clients send and receive secured e-mail
- Internal security policies have been established and clients are enrolled with respect to existing security policies
- Exchange Server provides mailbox and enrollment services
- Different S/MIME e-mail clients (Outlook 2000, Outlook Express, and non-S/MIME e-mail clients using Outlook Web Access) and enrollment processes (Key Management Service enrollment, Web-based enrollment, and so on) are used

  Roaming user support for certificates and keys for both internal and external users
- Certificate server hierarchy with root CA and issuing CA for generating and publishing certificates to Active Directory

Considerations for a Commercial Root CA vs. Self-hosted Root CA

A root CA does not have to be self-hosted. A commercial root CA can also be used. In this scenario, the Windows 2000 Certificate Services issuing CA is subordinate to the commercial root CA.

The advantages of using a commercial root CA include:

- Trust relationship is simplified, because all major commercial root CAs are trusted by default in Windows 2000 and in other operating systems
- Less administrative overhead in managing renewal and archiving root CA certificates
- Certificate Practice Statement is evaluated and insurance is provided in case of compromised root CA

The disadvantages of using a commercial root CA include:

- High level of trust accorded to external organization
- Strong dependency on external organization
- Organization's entire PKI affected if commercial root CA provider goes out of business

Building a Public Key Infrastructure

First, a PKI must be built, and a trust relationship for participating users. computers and services in the company must be established.

To build a PKI for the intra-company scenario

1. **Establish a root CA and distribute root CA certificates to participating entities** Group Policy in Windows 2000 manages the automatic distribution of root CA and issuing CA certificates. For other platforms or environments, a Microsoft Systems Management Server (SMS)–based distribution method or customized logon scripts may be necessary.
2. **Establish issuing CAs and intermediate CAs and build a CA hierarchy** In small companies, a hierarchy is not mandatory, but it is recommended. In large companies, a hierarchy is strongly recommended to address organizational, administrative, and scalability needs.

3. **Establish the necessary certificate templates and access policies** The use of certificates is determined by its attributes. Certificate Services in Windows 2000 supports predefined certificate templates, including attributes. Certificate templates are used to address different certificate functions and are secure.

4. **Plan internal client access and session types** The following clients are available for heterogeneous environments:

   - **S/MIME using Key Management Service** Key Management Service will automatically enroll users with minimal user intervention. This method requires that Key Management Service be installed and a Key Management Service strategy must be in place to address administrative and organizational tasks.

     Key Management Service S/MIME enrollment is supported only in Outlook 98 and Outlook 2000. Other SMIME mail clients use alternative enrollment techniques. Web-based certificate enrollment is a generic enrollment method for certificates.

   - **Non-S/MIME clients** For clients that are not running Outlook 98 or Outlook 2000, you can establish session-based security. A platform-independent option for accessing Exchange mailboxes is Outlook Web Access. To secure HTTP traffic, use SSL 3.0, a well-known security protocol. To initiate SSL, a Web server certificate is required. If you want to initiate a client authentication, a client certificate is required as well.

     Another session-based security protocol could be triggered with Microsoft MAPI clients, such as Outlook Express and the Exchange client. Access to mailboxes is handled using RPC. RPC supports encryption of RPC connections (secured RPC).

   - **Roaming S/MIME clients** These include internal users who roam within the company's domain and internal users who roam on computers that are not within a company domain.

     For internal users who roam within the company's domains, use the Windows 2000 roaming user profiles functionality. This makes SNIME key sets (certificates and private keys) available for the logged-on user.

     For internal roaming users who are not on computers within the company's domain, export the user's key set, which exists on the enrolled client, to disk and import the encrypted key set to the destination computer. Outlook 98 and Outlook 2000 support the .epf and Internet .pfx format. Outlook Express and other Internet e-mail clients support only the Internet Engineering Task Force (IETF) .pfx exchange format.

5. **Plan external client access** In this scenario, employees have a 1:1 relationship to customers. An SMIME client in the company is already enrolled. A major commercial CA (for example, VeriSign, Thwate, or GTE Cybertrust) provides the enrollment of the external users. See Figure 24.3 for a conceptual view of this scenario.

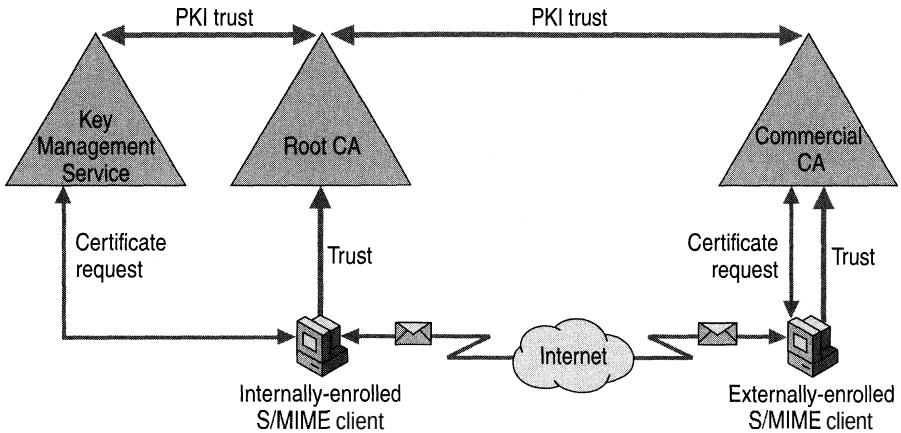


Figure 24.3 Trust for internal and external clients

For the scenario illustrated in Figure 24.3, you must build the PKI trust. By default, all major commercial CAs are trusted in Microsoft environments. The external user must trust your commercial CA. To establish this trust, your root CA certificate must be accessible to the external user.

Figure 24.3 describes how e-mail becomes secure for both the internally and externally enrolled clients. After the PKI has been established, the internal user sends a signed message to the external user. The external user retrieves the internal user's certificate and publishes this in the personal address book. The internal user does the same for the external user. Both the external and the internal users can now send and receive encrypted and signed messages at a peer level.

Secure communications are limited at the peer level. External users' e-mail addresses and certificates are available for just their one internal user. To make external users' information generally available for all internal users requires that you implement an ease-of-use process to self-publish external user's information in the internal company's e-mail directory. In addition, the external user has to be able to read and access the e-mail addresses and certificates of the internal company users.

Outlook 2000 Service Release 1 (SRI) addresses this self-publish consideration. With SR1, an enhanced Lightweight Directory Access Protocol (LDAP) client is provided, which allows retrieval of e-mail addresses and additional attributes (certificates) from an LDAP-based directory service, such as Active Directory. SR1 will also support a wizard for publishing certificates to Active Directory. Because this directory is accessible as an LDAP server, the

certificate is also accessible. There is currently no feature for publishing to arbitrary LDAP servers. In addition to online CRL checking, new features include security auto-configuration, support for Diffie-Hellman (a key Exchange algorithm), S/MIME version 3, and secure receipts.

Inter-Company Scenario

Many companies have business relationships with partners and curtomers. In some cases, the information that is exchanged between these partners is confidential and sensitive. Well-defined processes and a security infrastructure are required to ensure confidentiality and integrity of e-mail traffic between partners. Figure 24.4 shows PKI trust between a company running Windows 2000 Server and Exchange 2000 Server and company 2 running other systems.

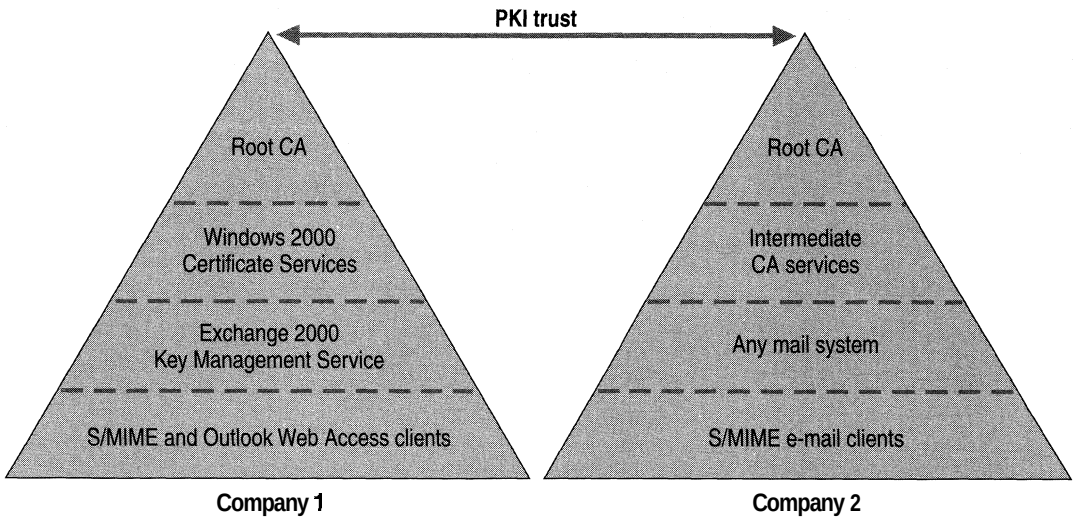


Figure 24.4 Inter-company security structure

Technical issues and features of this scenario include:

- Establish PKI at the company 2 site
- Exchange environment at the company 2 site is not necessary; however, the company 2 e-mail clients must support S/MIME
- User certificates and Simple Mail Transfer Protocol (SMTP) e-mail addresses must be accessible to both companies and company members
- Requires validation process and access rights strategy

Building a PKI Between Companies

Both partners in the inter-company scenario must first fulfill following prerequisites:

- Both companies must have an established PKI
- The partner's PKI does not need to be built on Windows 2000 Server, but it must support X.509 standards
- Both companies must agree upon and implement PKI trust relationships. In addition, both companies must agree on how to manage certificate management issues, such as CRLs, and accessing and exchanging certificates for S/MIME.

After the prerequisites have been fulfilled, a PKI trust between companies must be established.

- Companies must exchange their root CA certificates. Group Policy manages the automatic distribution of partner root CA certificates or certificate trust lists.
- Companies must define technical and organizational processes for the deployment of CRLs.
- Companies must implement a strategy that will allow users access to each companies' information, such as e-mail addresses and certificates, via LDAP. This is essential for exchanging secure e-mail.

Sharing Directory Information

Conceivable LDAP access strategies include:

- Direct access to directory service
- Directory service access using subdirectories (recommended)
- Directory service access using referrals
- Directory synchronization

Each of these strategies is discussed below.

Direct Access to Directory Services

In this model, both companies permit the partner company direct access to their own directory.

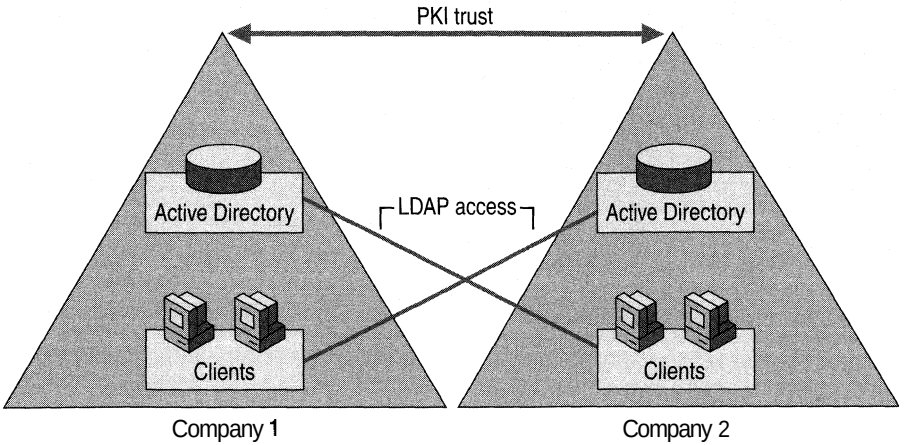


Figure 24.5 Direct directory access between companies

Each company's S/MIME client gets all relevant partner information, such as e-mail addresses and certificates to send encrypted and signed e-mail. An LDAP client, such as the Windows Address Book conducts the directory lookups. Another client option is the enhanced Outlook 2000 LDAP client that comes with Outlook 2000 SRI or Office 2000 SR1.

This scenario is rather improbable; for obvious security reasons, most companies do not want to provide direct access to their directories. In addition, this scenario requires a high administrative overhead to ensure access control to information that you do not want to share with a partner.

The implementation, however, is rather simple. LDAP and additional TCP ports must be opened to allow access *to* the directory service.

Subdirectory Access to Directory Services

Another method is to allow directory access via subdirectories. To do this, both companies establish a directory service for partner access; this directory contains a subset of the original company directory information. Only the necessary information for exchanging secure e-mail (e-mail addresses, user certificates, and user names) is published in this directory. Each participating company owns these subdirectories and is responsible for publishing their own data.

Figure 24.6 shows how the subdirectory is located in the perimeter zone (also called a demilitarized zone [DMZ]) in front of the firewall.

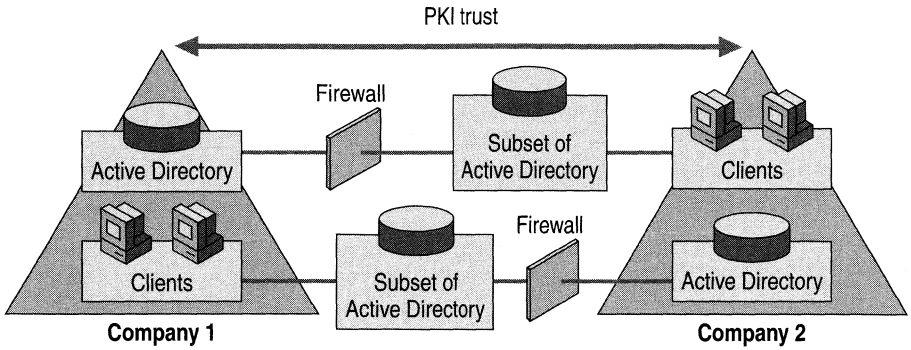


Figure 24.6 Subdirectory access between companies

This method is highly recommended, because few administrative tasks are required to establish appropriate access. In addition, only specific information can be accessed, while the majority of the directory is safe behind the firewall. Each company decides what to publish.

Referral Access to Directory Services

In this method, member companies browse for partner members in their own directory and are then referred to the partner's LDAP directory. The LDAP client receives referral information to generate an LDAP query. Although information retrieval is still through access at the partner site, this method enables central administration of the referral process—if the LDAP access point changes at the partner site, only the referral needs to be changed.

The partner must provide direct access to its directory, which has obvious negative impacts on security and administration. A combination of the referral scenario and the subdirectory scenario might be an interesting alternative to using only the subdirectory method.

Direct Directory Synchronization

In this scenario, the participating companies synchronize their directories. The companies must allow directory access to pull or push from each other's directories. Compared to the other scenarios, the administrative burden of this method is limited, and the access can be better controlled because access must be established only for a dedicated synchronizing service.

The LDAP client access to directory information resides in the boundaries of its own company, because all relevant information has been previously synchronized. One major issue exists: this scenario will work fine in an environment with a limited number of participating companies, but it can get hard to manage if several companies want to synchronize their directories with each other.

Inter-Company with Trusted Third Party Scenario

In some cases, the level of trust required for the inter-company scenario is not feasible. Sometimes, companies have loosely coupled relationships with partners, but also have demands to exchange secured e-mail. A disadvantage of the inter-company scenario is the high administrative burden required to set up such a trusted and secure environment.

Establishing a trusted third party will help alleviate some of the administrative burden inherent in the inter-company scenario. In addition, trusting a third party does not require such a high degree of trust to be given to a single partner. A trusted third party can be an independent organization, such as VeriSign, or an arrangement can be made by the company to set up its own trusted third party. The success of this scenario depends on a scheme in which all participating companies and partners trust the third party CA.

A security infrastructure for a closed group can be established using the third party without any direct trust relationship between partners. A trusted third party addresses the challenges of creating a highly secured e-mail environment for loosely coupled relationships, and provides a way to decrease the administrative burden of managing the CA.

The following list discusses some of the technical issues and features of this scenario:

- The security infrastructure in the company and at the third party site is similar to the inter-company scenario.
- PKI Trusts must be established and maintained by the third party CA.
- An LDAP-based directory must be built at the third party's site to store all users' e-mail addresses and certificates.
- Access policies and network security must be defined and maintained by third party.
- The distribution process of certificate trust lists, CRLs, and other certificates that are exchanged between companies must be defined.

Building a PKI With a Third Party

Setting up a PKI with a third party will require some extra planning up front. The following questions must be considered as you plan your third party CA strategy:

- Who will be responsible for the trusted third party?
- Who owns the third party?
- Who will be trusted for this service?

Figure 24.7 shows the PKI trust between two companies and a trusted third party. This model can extend to many companies.

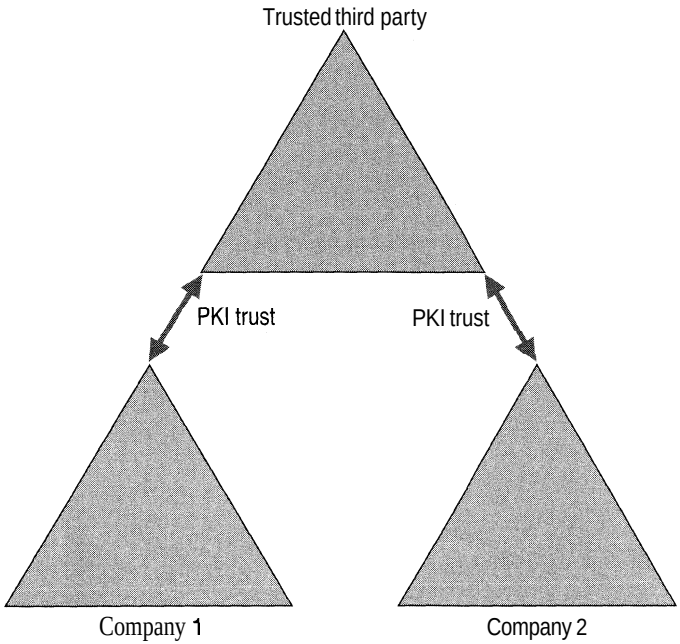


Figure 24.7 Inter-company security structure with trusted third party

Building a PKI trust relationship with a trusted third party can be approached in different ways. The first major considerations are who will sign the certificate trust list and where the certificate trust list will be generated.

The decisions you make regarding the certificate trust list depend on organizational issues. The level of trust (distinguished in this discussion as ***very*** *high*, or *high)* granted to the third party is determined by the company. The trust is considered very high for a self-owned third party and high for an external commercial third party.

The following two scenarios are possible:

- Trusted third party with very high trust level
- Trusted third party with high trust level and directory service available over the Internet

Very High Trust Level

In a very high trust level environment, the main administrative tasks, which include establishing a PKI trust between companies and their partners, are delegated from the company site to the third party.

First, a trust between the company and trusted third party must be established. To do this, root CA certificates must be exchanged between the company and trusted third party.

The next step involves the PKI trust relationship between companies. The trusted third party builds the company 1 and company 2 certificate trust list and signs the certificate trust list with its own Trust List Signing certificate issued by the trusted third party's CA. The participating companies retrieve the certificate trust list generated by the trusted third party. Companies can verify the certificate trust list signature as valid, because the PKT trust to the trusted third party exists.

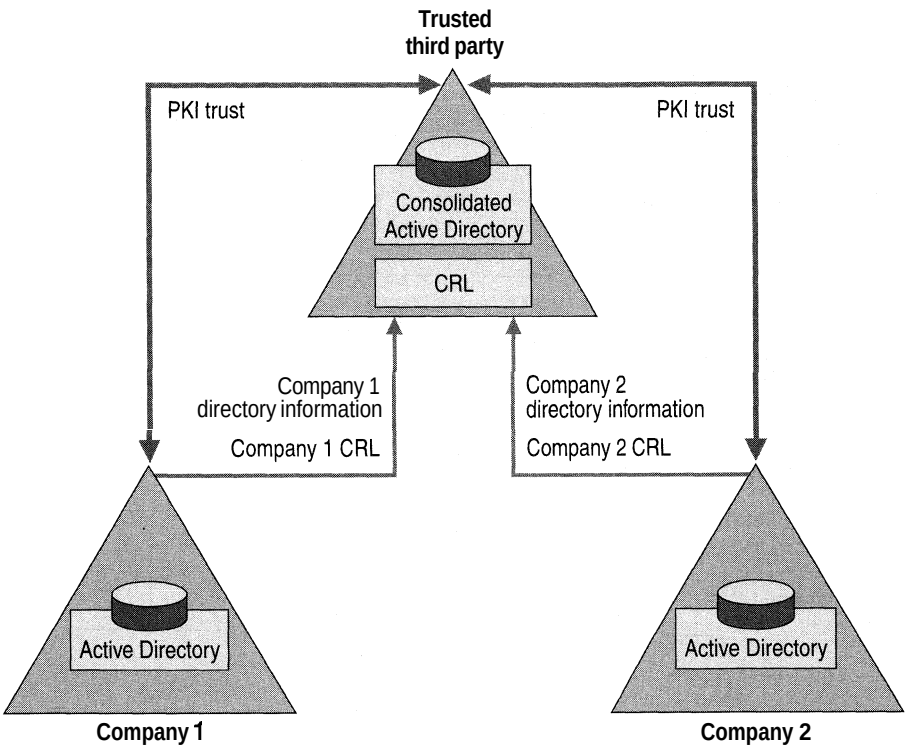


Figure 24.8 Trusted third party with very high trust

Internal certificate trust list distribution is similar to the root CA certificate distribution using Group Policy, which is discussed in the preceding section, "Intra-Company Scenario."

There is a possible variation to this scenario. The trusted third party-generated certificate trust list can also be signed with the trust list signing certificate issued by each company. However, this can be difficult to manage for several companies because the same certificate trust list has to be signed with all Trust List Signing certificates to validate the certificate trust list at the company site. In the simplest case, if all companies (assuming that there are three or more companies) trust each other, the trusted third party distributes one certificate trust list signed with different certificates. However, if Company 1 trusts Company 2 but not Company 3, separate certificate trust lists are required and the complexity increases enormously.

High Trust Level

In this scenario, the trusted third party is a commercial service hired by participating companies This provides for a high trust level environment in which each company performs the administrative tasks for establishing a PKI trust. Thus, building a PKI trust to the trusted third party is not mandatory. However, if your trusted third party provides optional security features, such as SSL using HTTP, LDAP, or Internet Protocol security (IPSec), you can establish a PKI trust between your company and the trusted third party.

Figure 24.9 shows each company publishing its relevant directory information to a consolidated Active Directory owned by the trusted third party. Each company also publishes its root certificate and CRL. The trusted third party makes the companies' root certificates available and allows each company to build certificate trust lists (from the consolidated directory).

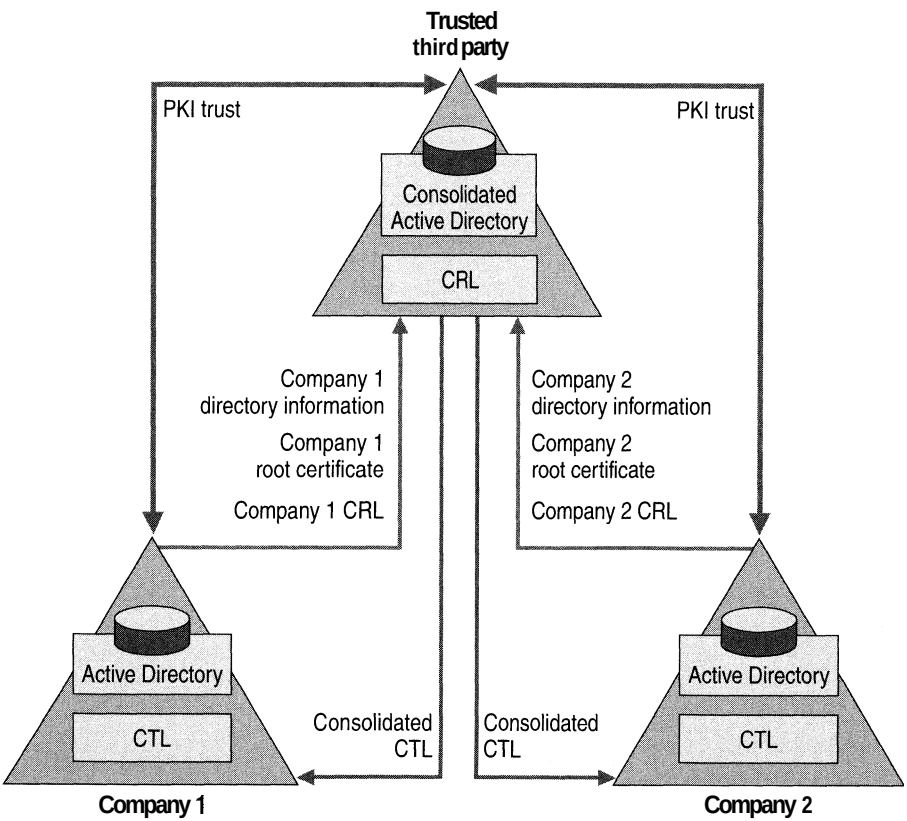


Figure 24.9 Trusted third party with high trust

A PKI trust is not needed when referring to commercial CA certificates (for example VeriSign) in the following scenarios:

- The participating companies publish their own root CA to a distribution point located at the trusted third party. The trusted third party provides access to the companies' root CA certificates.
- Each participating company retrieves the appropriate root CA certificates and builds its own certificate trust lists signed with its own Trust List Signing certificate.

CRL and Directory Access

To enable access to CRLs and LDAP directory information the following tasks are common for both high trust level and very high trust level scenarios.

All participating companies determine two CRL distribution points. One CRL distribution point is located in the company site for internal users to check for revoked certificates, and a second CRL distribution point is located in the trusted third party site for the external users. The second CRL distribution point provides a revocation checkpoint for the company certificates.

In other words, company users check company certificates against the CRL published at the company CRL distribution point, and partner users check company certificates against the CRL published in the trusted third party site.

Therefore, the trusted third party must establish a CRL distribution point for all participating companies. Preferably this would be an HTTP-based distribution point.

Each participating company pushes their CRLs to the trusted third party-hosted certificate distribution point.

> **Note** Certificates must support the CRL distribution point attribute. The CRL distribution points are listed in the certificate itself. CRL distribution points are checked in sequential order. The first positive CRL check for a revoked certificate will stop all following CRL checks from other CRL distribution points. Certificates that are already enrolled will not check against the trusted third party-hosted CRL distribution point.

The trusted third party must provide access via LDAP or via directory synchronization to all participating companies users' e-mail addresses and certificates. The scenarios describing how to access and publish directory information are discussed in the previous section, "Inter-Company Scenarios."

To enforce secure access to trusted third party's hosted directory, several security mechanisms can be used. You can use LDAP to retrieve directory information at the trusted third party site, and you can secure HTTP access for hosted CRLs with SSL. You can use IPSec to secure publishing or synchronizing traffic. And finally, a firewall can control all network traffic to and from the trusted third party.

Deployment Scenarios

The purpose of the following section is to provide a practical guide for configuring the previous Key Management Service scenarios. Although the configurations that follow were installed and tested in a lab environment, they will help you understand some of the issues to consider when planning the deployment process of an infrastructure supporting secure messaging with Exchange 2000. That said, this section is not intended to be a walkthrough that breaks down deployment into a detailed step-by-step process. For procedural specifics regarding PKI, refer to the Exchange 2000 Server documentation.

Intra-Company E-mail Security

The first scenario describes the basic implementation of a messaging environment for one organization that implements a PKI infrastructure for S/MIME.

Before starting with the deployment of the CA infrastructure the following issues should be considered:

- Does the High Encryption Pack for Windows 2000 need to be installed to support the Microsoft Enhanced Cryptographic Provider? If high encryption is needed, the deployment must include the Outlook High Encryption Pack.
- Check other platforms and applications to verify their compatibility for high encryption. If necessary, update the platforms and applications to the appropriate versions.

  Note Although the High Encryption Pack for Windows 2000 is only needed on the CA servers, consider deploying the High Encryption Pack on Exchange 2000 Key Management Service servers, and on the Outlook clients. This will ensure a consistent high encryption environment. It doesn't make sense to run the High Encryption Pack only on the CA servers. For example, consider SSL connections to Web servers that do not have the High Encryption Pack installed; this is effectively downgrading clients with high encryption to normal encryption.

- How many levels (intermediate and issuing) of CAs are needed to support the requirements of the organization?
- Be prepared to spend more time installing the root CA than subsequent CAs. Typically, you install the root CA with a long expiration period (maybe 5 years) and a public key digital signature length of 2048 bytes. The root certificate is exported for signing additional CAs, and the computer is taken offline and placed in a secure area.
- Familiarize yourself with the X.509 ITU specification about the format and naming conventions for the attributes for CAs.

The following intra-company deployment scenario will run high encryption and will establish a two-level CA hierarchy.

Intra-Company E-mail Security Scenario

Before starting deployment of the PKI, install the High Encryption Pack for Windows 2000 on every Windows 2000–based server and on every Windows 2000–based client. Also install the High Encryption Pack for Outlook 2000 on all clients running Outlook 2000.

The next step is to install Certificate Services as an enterprise root CA supporting the enhanced cryptographic service provider with a public key length of 2048 bytes. Then, install an enterprise subordinate CA to issue certificates to users and computers. Check that the enterprise root CA Certificate is in the Trusted Root Certification Authority Store of the computer. This is an automatic process enforced through Windows 2000.

> **Note** Typically, issuing CAs are installed with a shorter public key due to scalability considerations.

Before installing the Exchange 2000 Key Management Service, add the necessary templates: **Exchange User. Enrollment Agent (Computer),** and **Exchange Signature Only.** The computer account where the Key Management Service will reside is given rights on the issuing CA to Manage, Enroll, and Read. Again, the enterprise root CA Certificate needs to be installed on the Key Management Service computer.

Installing the Key Management Service is now a straightforward process. Choose the startup method, which defines how the password for the Key Management Service is maintained, and determine which Exchange 2000 administrative group this Key Management Service is responsible for.

> **Note** It is absolutely necessary to keep the password for the Key Management Service in a highly secured place to prevent intruders from breaking into the system and recovering keys.

The existing environment is illustrated as part of figure 24.10. It consists of one Active Directory domain with one domain controller, the root CA, one issuing CA and one server running Exchange 2000 Server and Key Management Service (Exchl). All servers are running Windows 2000 Advanced Server. Clients are running either Windows 2000 Professional with Outlook 2000, Outlook Express, or Internet Explorer 5; or a UNIX system with Netscape Communicator 4.7.

Enrolling S/MIME Certificates

Key Management Service provides additional enrollment features to Outlook 2000 and Outlook 98 clients for S/MIME. Through Exchange 2000 System Manager, the enrollment process can be initiated directly through the Key Management server. You can enroll users by selecting users from an address list that consists of users with mailbox resources on servers in a specific administrative group, or by selecting mailbox stores available in a specific administrative group.

Enrolling certificates for Outlook Express and Netscape Communicator is done through the Certificate Web Enrollment tool on the issuing CA.

To further configure Outlook Express, you can set the security properties for the Signing and Encryption Certificate through the mail account properties dialog box. Outlook Express can be configured to use a dual key pair by enrolling two certificates: User Certificate and User Signing Certificate.

Accessing the Certificates for S/MIME

All certificates enrolled by Certificate Services (and therefore by Key Management Server in S/MIME mode) are stored in Active Directory. Outlook 2000 uses Active Directory as the global address list so all certificates are automatically available to Outlook 2000 clients. Other clients can access Active Directory through LDAP, getting the necessary certificates to encrypt e-mail. One such LDAP tool is the Windows Address Book that works together with Outlook Express. For users running Outlook Express on Windows 2000 Professional, the default entry, Active Directory, already exists so that clients can look up directory information in Active Directory.

Enrolling SSL Certificates

Running Outlook Web Access over SSL requires a Web Server Certificate for Internet Information Services (IIS). This is done through the Web Server Certificate Wizard for the Web site hosting Outlook Web Access. After the Web Server Certificate has been installed and configured, SSL needs to be enabled through the properties on the Exchange virtual root. In addition, optional 128-bit enforcement for high encryption can be configured.

To work with client authentication, a user certificate must be enrolled on the client. Mapping the certificate to a user in Active Directory is then an easy step: In Internet Services Manager, right click the server name in which IIS is running and click **Properties.** On the **Internet Information Service** tab, in **Master Properties** click **Edit.** On the **Directory Security** tab, check **Enable the Windows directory service mapper.** This allows clients that can only work with basic and clear text authentication to securely connect to Outlook Web Access by authenticating with their certificate (mutual authentication).

Intra-Company E-mail Security Scenario

The second scenario adds another administrative group to the Exchange organization with another Key Management server. As illustrated in figure 24.10, the environment consists of one Active Directory domain with one domain controller, the root CA, one issuing CA and two servers running Exchange 2000 Server and Key Management Service (Exch1 and Exch2). The same considerations apply as in scenario one. Both Key Management servers use the same issuing CA.

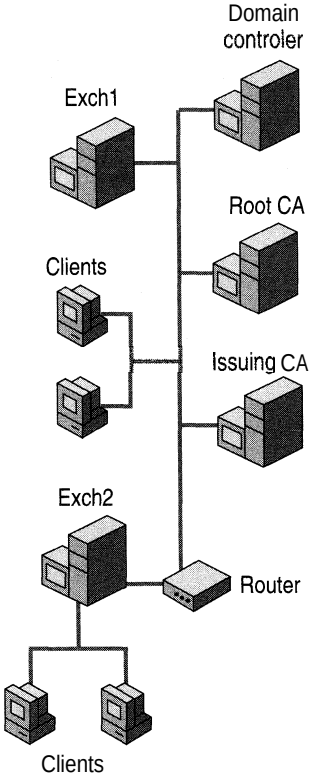


Figure 24.10 Intra-company e-mail security

Exchange 2000 Key Management Service provides great flexibility for operating with issuing CAs. It actually gets any available CA and uses the one that is returned. To force Key Management Service to use a specific issuing CA, take the other CAs offline, establish a connection from Key Management Service to the specific issuing CA, and then bring the other issuing CAs online. If your deployment strategy requires specific CA mapping, take this into account in your design planning considerations. In addition, possible fail-over and fail-back issues should be accommodated by your plan.

Intra-Company E-mail Security Scenario

Depending on your organizational structure and your administrative needs, the previous scenario might not apply to you.

In the intra-company e-mail security scenario, the S/MIME e-mail security for the second Exchange administrative group is established with an additional issuing CA and Key Management server.

This is a scenario that can be implemented in situations where independent divisions want to have the authority for their PKI.

Note here, that although it is not necessary for the S/MIME e-mail security, Windows 2000 is configured through policies to allow communication between sites using only IPSec.

Keep in mind that if such security requirements exist, communicate with the appropriate Information Technology staff to ensure that required ports and filters are configured according to appropriate security policies.

Also be aware that encryption may disable additional security features like virus checking, which is the primary function of some firewall implementations.

Inter-(E-mail it

The scenario described in this section explains the requirements for enabling two separate companies to communicate using S/MIME.

The following needs to be in place or planned before implementing the S/MIME infrastructure in this scenario:

- Both companies must have a PKI in place.
- Both companies must agree upon and implement PKI trust relationships. In addition, both companies must agree on how to manage certificate management issues, such as CRLs, and accessing and exchanging certificates for S/MIME.

Installing a PKI with Key Management Service for S/MIME is discussed in the Intra-Company Scenarios section. The concepts explained previously can be applied to both companies separately.

The following additional steps must be followed:

- Defining the access method to the certificates of the companies.
- Defining how the root certificates of the companies are imported.
- Defining how CRLs will be managed.

In this scenario, configure directory replication between both directories to synchronize user objects from one company to import as contacts to the other company.

CRLs of one company are copied to a Web site of the other company for end users to download. A more sophisticated solution will be discussed in the next scenario and can be also applied to this solution.

> **Note** The CRL is a file stored by default in the //certsrv directory with the name of the CA and the extension .crl.

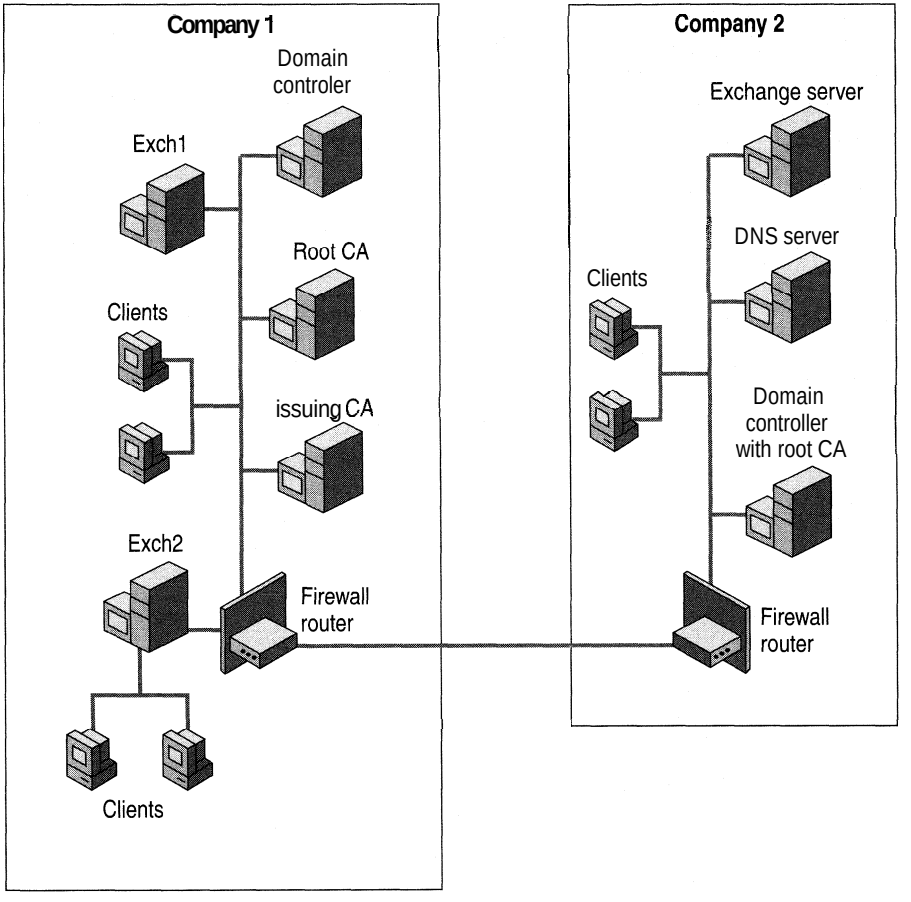


Figure 24.11 Inter-company e-mail security

Figure 24.11 shows the environment of the Inter-company scenario. Note again that a *firewall* is configured to allow for a tightly controlled communication between companies.

Because the directory of each organization is available in Active Directory, no additional configurations are necessary for the clients to access certificate information.

Directory Replication Between Two Forests

At the time of this writing, direct synchronization of Active Directory forests is not yet available in Windows 2000. However, because many organizations use Exchange 5.5, using Exchange 5.5 solely as a directory store and configuring the Windows 2000 Active Directory Connector (ADC) between two forests can solve this problem.

This is not a perfect solution, especially when more than two companies are involved in a direct PKI relationship. Configuring directory replication across companies using ADC needs to be planned and deployed carefully to avoid administrative problems. Evaluate products that are focused on directory synchronization, like Microsoft Meta Directory Services to implement Inter-forest synchronization.

For more information about Active Directory Inter-Forest Synchronization, see "Inter-Organization Replication and Directory Synchronization" in this book.

Smart Cards

In this scenario you can introduce smart cards to test how messaging clients work with certificates generated by a smart card cryptographic service provider. Keep in mind that available smart cards allow only for single key pair encryption.

As the smart card is itself a cryptographic service provider, the first challenge is using smart cards in combination with Key Management Server. Key Management Service cannot use the cryptographic service provider on the smart card to generate the encryption certificate. Another thing to consider is that Outlook uses its own implementation to generate the signing certificate in conjunction with the Key Management Server enrollment. Given these issues, Key Management Server cannot be used for smart card users.

However, it is possible to use smart cards for S/MIME security via the normal enrollment process available with Certificate Services in Windows 2000. Using the certificate template for a smart card user, a certificate is generated that can be used not only for the Windows 2000 logon, but also for S/MIME. In this scenario, you will need to configure the security options in Outlook 2000 and Outlook Express accordingly. Certificates enrolled in this way for smart cards can be used without any further preparations for SSL-based security access, such as Outlook Web Access.

When you send e-mail from a dual key pair-enrolled Outlook 2000 user to a single key pair-enrolled user, Outlook 2000 does not recognize the enrolled certificate for both (dual key) purposes and accepts the certificate, but not as an encryption certificate.

> **Note** Typically, smart card enrollment follows the same methodology as Key Management server does. Smart card certificates are generated on behalf of a user. The prepared smart card is then sent directly to the user. To enroll certificates on behalf of another person, the person who executes the enrollment needs to have an Enroll Request Agent (User) certificate.

Trusted Third Party (Internet Directory Services)

This scenario shows an implementation with a trusted third party that acts as a directory service for separate companies that need to communicate securely.

This trusted third party scenario requires that every participating company has established a PKI. Introducing a trusted third party streamlines many PKI processes that otherwise need to be organized independently from each company. This solution could be the vision for every inter-company PKI, as it also streamlines the management processes for scalability and flexibility.

Figure 24.12 illustrates a trusted third party infrastructure. Compared to the Inter-company scenario, another Active Directory forest is established to hold the trusted third party's directory.

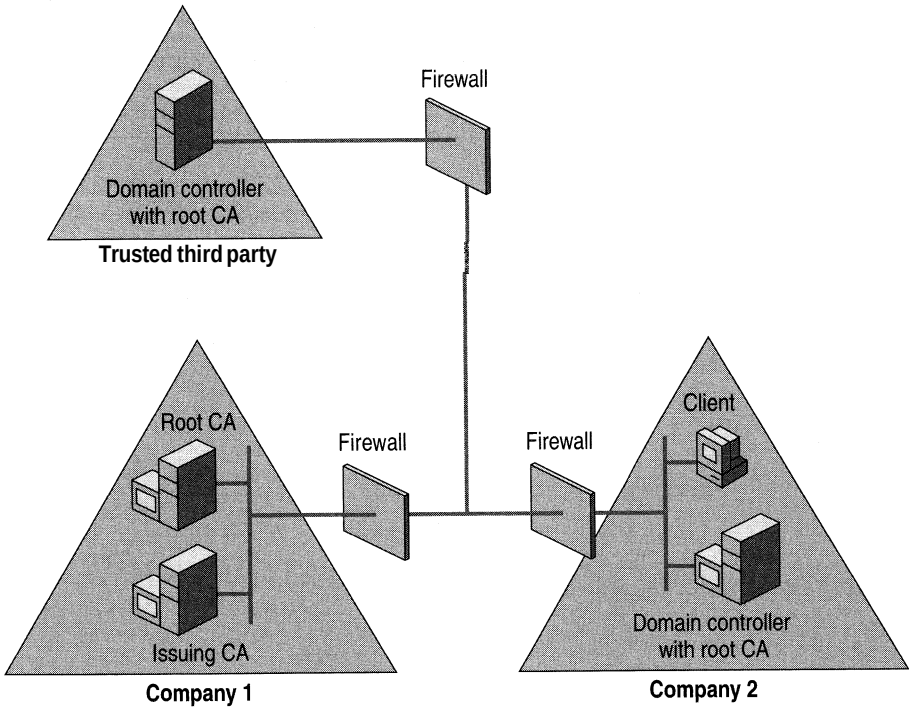


Figure 24.12 Trusted third party scenario

Planning considerations include defining the access method to the certificates, defining how the certificate trust lists are managed and imported to the organizations, and defining how CRLs will be managed.

The scenario, as discussed here, was implemented in the lab environment in the following ways:

- By replicating all directory information from both organizations to the trusted third party
- By enabling the clients to access the trusted third party, not importing directory information from other companies
- By building certificate trust lists on each organization separately
- By publishing CRLs from each organization to a distribution point on the trusted third party

Replicating Directory Information to a Trusted Third Party

As in the Inter-company scenario, Microsoft Exchange 5.5 and the Windows 2000 Active Directory Connector (ADC) can be used to import user objects with mailboxes as contacts with certificates to an Active Directory directory service maintained by the trusted third party. Note that this must be planned carefully and that a direct directory access might not be appropriate for most scenarios.

Accessing Directory Information in a Trusted Third Party

In this scenario, directory information from the companies is not directly imported into each other's directory. For this reason, Outlook 2000 and the Windows Address Book for Outlook Express have to be customized to access encryption certificates. All the objects with the required information are stored in the trusted third party directory.

The first step in this customization is to enable anonymous LDAP access on the Active Directory directory service of the trusted third party. Plan for additional security so that Active Directory is not open for everyone. Using SSL for LDAP or IPSec could be a solution.

The next step is to configure Windows Address Book by adding another account for the trusted third party directory. The access can be set on the global catalog LDAP port 3268.

The same is done in Outlook 2000 with the Microsoft LDAP Directory provider.

> **Note** To retrieve certificates with the LDAP provider in Outlook 2000, you must be running Outlook 2000 SR1 or Office 2000 SR1.

Building Certificate Trust Lists for Each Company

The certificate trust list in this scenario is built separately for each company.

Each company must publish its root CA to the trusted third party. Each company must then download the root CA for the trusted third party and the other company's root CA certificate. The certificate trust list is then built from the downloaded root CA certificates.

To import the certificate trust list, the account used to install the certificate trust list must have a Trust List Signing Certificate. This is used to sign the certificate trust list prior to import. To distribute the certificate trust list to all computers in the Active Directory domain, the certificate trust list can be imported through the Default Domain Group Policy (other Group Policy assignments may be used on different levels).

Publishing Certificate Revocation Lists to a Trusted Third Party

CRLs for this scenario are accessed through distribution points configured as virtual roots on an IIS server of the trusted third party.

Each company publishes its CRL through the Certificate Authority Administration to a file that is uploaded to the configured distribution point.

> **Note** The certificate distribution point in the certificates must include the internal path and the path of the published CRL at the trusted third party. In this scenario, the certificate distribution point is adjusted by adding the virtual root to the file name on the trusted third party. This is done through the Certificate Authority Administration properties.
>
> Keep in mind that certificates that have already been enrolled do not include the modified certificate distribution points.

Next, clients must be configured to automatically check against the certificate distribution point in the certificates. Outlook Express and Outlook 2000 allow for this functionality. Outlook Express exposes this in the **Advanced Security** tab of the **Options** dialog box. Outlook 2000 doesn't expose this feature, but it's available by editing the registry. Aside from the registry edit that enables the CRL lookup, there's a second registry edit that defines a severity flag, which determines how Outlook should react if the CRL is not accessible.

You can configure the registry setting that determines whether a missing CRL for a signature is an error or a warning. If you set the flag to be an error, the signature will not be validated. By default, the flag is set to a warning, so that the signature is accepted and the user receives a warning message.

To set the behavior for a missing CRL

> **Caution** Do not use a registry editor to edit the registry unless you have no alternative. Registry editors bypass the standard safeguards provided by administrative tools. These safeguards prevent you from entering conflicting settings or settings that are likely to degrade performance or damage your system. Editing the registry directly can have serious, unexpected consequences that can prevent the system from starting and require that you reinstall Exchange 2000 Server or Windows 2000 Server. To configure or customize Windows 2000, use the programs in Control Panel or Microsoft Management Console (MMC) whenever possible.

To enable CRL checking in Outlook 2000

1. On the **Run** line, type **regedt32.exe** or **regedit.exe**, and then click **OK.**
2. In the registry editor, navigate to HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Cryptography.
3. Add or modify the {7801ebd0-cf4b-11d0-851f-0060979387ea} subkey.
4. To enable CRL checking add a **PolicyFlags** entry and assign data type **REG–DWORD** and value **0x00010000.**
5. Close the registry editor.

To set the behavior of CRL checking

1. In a registry editor, navigate to HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Office\9.0\Outlook\Security.
2. Click the **SigStatusNoCRL** entry.
3. In Regedit.exe, right-click the entry, and then click **Modify.**

   –or–

   In Regedt32.exe, click the entry, click **Edit,** and then click the appropriate menu choice.
4. To set the setting so a missing CRL generates a warning, assign a value of **0**.

   To set the setting so a missing CRL generates an error, assign a value of **1.**
5. Close the registry editor.

   Caution CRLs can have a side effect when the CRL distribution point is not available. For example, enabling CRLs in Outlook 2000 and working offline will generate long waiting times because Outlook will try to access the CRL distribution point. This behavior should be expected for all applications using advanced security with CRLs.

   Note Even if you turn on CRL checking, Outlook will only look for new CRLs if the existing cached one expires.

Internet Explorer also exposes a setting in the **Internet Options** and **Advanced** menu to enable online checking of CRLs.

CRL checking can also be delegated to the operating system. By importing CRLs to the local computer and registering the CRL in the local certificate store (with scripting or deployment tools like SMS), Windows 2000 will first check the local certificate store and will find revoked certificates.

In general, CRL checking is accomplished by first checking the local certificate store, and then sequentially checking all CRL distribution points in the certificate. If a matching revoked certificate is found, checking ceases.

CHAPTER 25

Outlook Web Access

Bob Hunt, Managing Consultant, E-Sync Networks, Inc.
Carl Solazzo, Enterprise Consultant, E-Sync Networks, Inc.
Paul Sebben, Managing Consultant, E-Sync Networks, Inc.

Microsoft Exchange 2000 Server introduces several fundamental architectural changes that improve the functionality and scalability of Outlook Web Access. Figure 25.1 illustrates the new architecture of Outlook Web Access.

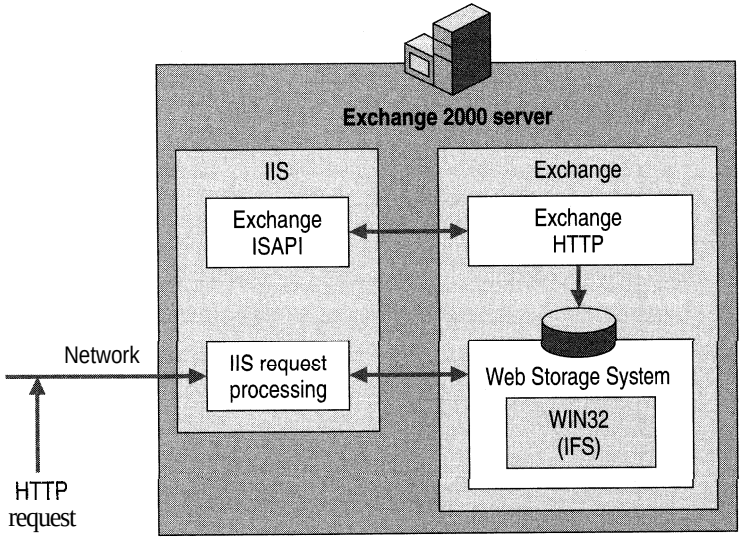


Figure 25.1 Outlook Web Access architecture

Microsoft Internet Information Services (IIS) is now coupled with Exchange, and is required on all Microsoft Exchange 2000 servers. IIS is now responsible for client protocol support and it manages all client requests for the following Internet protocols: Hypertext Transfer Protocol (HTTP) with Web Distributed Authoring and Versioning (WebDAV) support, Post Office Protocol version 3 (POP3), Internet Message Access Protocol version 4 (IMAP4), Network News Transfer Protocol (NNTP), and Simple Mail Transfer Protocol (SMTP).

The Microsoft Web Storage System includes native HTTP access and manages data for Exchange. Every object within the Web Storage System is URL–accessible with short, easily understood names. URL escaping and encoding is necessary for non-ASCII characters. The Exchange 2000 Installable File System (ExIFS) is optimized for streaming large Multipurpose Internet Mail Extensions (MIME) data. Because the Web Storage System stores data in its native format, no data conversion is necessary, and therefore data can be quickly retrieved. HTML rendering is no longer script–based.

In This Chapter

Outlook Web Access Evolution

Features

Clients

Deployment Planning

Front-End and Back-End Server Architecture

Authentication

Performance Monitoring

Outlook Web Access Evolution

Microsoft introduced Outlook Web Access as an optional component in Microsoft Exchange version 5.0. Outlook Web Access is a viable solution for companies for the following reasons.

- Users can gain access to their mailboxes from virtually anywhere within the company intranet, or even from the Internet—without needing to reconfigure the client profile or install additional software. For example, one shared client system with a Web browser can service users for an entire factory floor.
- Outlook Web Access provides users on UNIX and other operating systems with access to their e-mail by using a Web browser.
- Deployment of an e-mail client is not required. Companies can support telecommuters without needing to deploy the e-mail client on their home system.

In Microsoft Exchange 5.5, the Outlook Web Access server performs most of the processing for the connected clients. Outlook Web Access for Exchange 5.5 was built using interpreted Active Server Pages (ASP) scripts. For each connected user, Outlook Web Access must open a MAPI session with the Exchange server, but the number of open MAPI sessions that can be efficiently handled within an ASP session is limited. The architecture restricts the number of users that can simultaneously connect to an Outlook Web Access server. The number of MAPI sessions does not limit the number of users you can host with Outlook Web Access. In fact, when accessing a typical Exchange server by using Outlook, you have the same number of MAPI sessions. The problem is that Outlook Web Access for Exchange 5.5 was comprised of interpreted ASP script that could not accommodate a large number of open MAPI sessions.

Features

The following sections introduce the most significant changes for Outlook Web Access in Exchange 2000 Server.

Accessing Your Mailbox

Because every object within the store is URL–accessible, users have several different ways to access objects within mailboxes or public folder hierarchies. The URL for an object is based on its location within the hierarchy and generally contains the subject of the item. If you are using Microsoft Internet Explorer and have already authenticated in an Active Directory directory service domain that is trusted by the domain hosting Exchange 2000 Server, no authentication prompt appears.

If you are not authenticated, the Web browser prompts for a logon ID. domain, and password. Also, if you access an Outlook Web Access server that does not contain your mailbox, Outlook Web Access performs a search for the appropriate server and redirects the client session to that server. The following are some examples of accessing Exchange resources using Outlook Web Access:

To open a mailbox

- Type **http:llsewernamelexchangeluseridl**or http:llsewernamelexchangel

To open the top-level public folder

- Type **http:llservernamelpublicl**

To open your calendar

- Type **http://*servername*/exchange/*userid*/calendar/**

Sewername is the name of the Exchange server and exchange is the virtual directory that points to your users' private folders. *Userid* is taken from the left half of the user's SMTP address. That is, if a user's SMTP *address* is bob@winery-co.com the URL to access Bob's mailbox will be http:llservername/exchange/bob.

For example, if the *exchange* virtual directory points to m:\winery-co.com\MBX, then users must have an SMTP address of bob@winery-co.com to log on to Outlook Web Access. In some cases, it is necessary to add appropriate secondary SMTP addresses to users that are moved or upgraded to enable them to log on to Outlook Web Access.

To exit the Outlook Web Access session, close all instances of the browser. This clears any authentications that may be in the browser cache.

Accessing Exchange Objects

This section describes the sequence that occurs when a user accesses an object within Outlook Web Access. Figure 25.1 illustrates the architecture described here.

The mailbox is opened, and a user clicks a message to open it. The IIS request processor calls the Exchange HTTP Internet Server Application Programming Interface (ISAPI) application that parses the information in the request and determines:

- The action to be performed (open mailbox, open folder, read mail, create mail, and so on)
- Browser information (browser type, version, how the information should be rendered)

The server then determines whether the user has rights to gain access to the item. If so, then the object state (read, unread), object type (folder, message, and so on), and item type (message, appointment, contact) are determined.

The Exchange HTTP ISAPI extension then matches the object attribute to its corresponding form definition. If a form definition does not exist for a particular object attribute, then the default form is used, (the one used to read an e-mail item). The Exchange HTTP ISAPI extension then parses the form and queries Web Storage System to bind to the data. After receiving the data from the Web Storage System, the Exchange HTTP ISAPI extension renders the data in HTML or Extensible Markup Language (XML) based upon the browser type and version. The client can then view the message.

WebDAV and XML

Web Distributed Authoring and Versioning (WebDAV) is an extension to the HTTP 1.1 protocol (RFC 2518). HTTP and WebDAV allow rich collaborative interaction with the Web Storage System in Exchange 2000. Exchange 2000 HTTP support allows for adding, modifying, copying, moving, and searching of folders, items, and manipulation of attributes on any object in Web Storage System.

In the context of Outlook Web Access, WebDAV creates improved performance and user experience over the HTML 3.2 client by exploiting client-side data binding and rendering. For example, when you click on the column header, you can sort the Inbox in several different ways, allowing views based on the sender's name, the message subject line, or received date. The browser caches the user interface (UI) elements such as Internet Explorer 5.0 HTML Components, Microsoft JScript libraries, Extensible Stylesheet Language (XSL), and Graphics Interchange Format (GIF) files. When the user changes the sort criteria, the browser can reformat the UI elements locally and query for the view data from the server.

Clients

Outlook Web Access supports several different Web browser versions including Netscape Navigator 4.08 and higher, Internet Explorer 4.01 Service Pack 1 (SP1), and Internet Explorer 5.01 and higher. When the Outlook Web Access session starts, the server detects the browser type and version and provides the appropriate user environments.

Clients that use Outlook Web Access support calendar functions and user contacts within the public folders, address card views, and the ability to resolve names against personal contacts. Columns within the mailbox can be sorted when the header line is clicked. Outlook Web Access now supports reading embedded objects inside of messages.

Internet Explorer 5.0

Internet Explorer 5.0 is the preferred Web client (also called rich client) for Outlook Web Access in Exchange 2000, and provides a user experience similar to Outlook. This client supports the latest Internet technologies such as WebDAV and XML. Currently, only Internet Explorer 5.0 running on Microsoft Windows 95, Microsoft Windows 98, Microsoft Windows NT 4.0, and Windows 2000 supports these technologies. New technology that Outlook Web Access supports includes:

- **Preview pane** Allows you to preview a message prior to opening it.
- **Drag and drop functionality** Allows you to manage the messages within your mailbox by providing the ability to move documents between folders within the mailbox.
- **HTML text editing** Provides you the ability to change font size, style, and color within the browser when sending or replying to messages.
- **Right-click menu options** Provides easy management in the folder hierarchy.

In addition to the functionality enhancements, Internet Explorer 5.0 also offers enhanced performance. With the technology available with Internet Explorer 5.0, the Web browser processes many Outlook Web Access commands without sending a request to the server. This reduces network traffic between the server and the browser.

Other Clients

Due to the popularity of the Internet, Outlook Web Access can make Exchange e-mail accessible from many public kiosks in airports, restaurants, and other commercial buildings. In these locations, users must use the available Web browser.

Some companies may not have the ability to deploy Internet Explorer 5.0 in a timely fashion. Therefore, it is necessary that a Web client be available to support other browsers (also called reach clients). These include browsers—other than Internet Explorer 5.0—that comply with the HTML 3.2 and European Computer Manufacturers Association (ECMA) script standards.

Support for other browsers is greatly improved since earlier versions of Exchange. No add–ons are required to access Outlook Web Access in Exchange 2000, which reduces the compatibility issues in earlier versions of Outlook Web Access, and increases performance. Support for other browsers includes two frames. The left frame is the navigation frame and is static. The right pane displays folder views, messages, appointments, contacts, and so on.

Client Limitations

Despite improvements, Outlook Web Access in Exchange 2000 still has limitations in comparison to Outlook 2000. Outlook Web Access does not support working offline, spell checking, the ability to create Outlook rules, or to gain access to tasks and journal items. Messages cannot be sent for deferred delivery, or set to expire, and you cannot copy messages between public folders and mailbox folders. Also, you cannot create or manipulate digitally signed or encrypted messages, although these limitations are due to limitations in the current Web browser architecture.

Deployment Planning

Outlook Web Access is automatically installed when you install Exchange 2000. When the number of people who use the server grows, IIS must process more protocol requests on HTTP. The server that contains mailbox stores is called a back-end server (a server that houses mailbox stores and public folder stores). The role of a front-end server is to communicate directly with the clients' browsers and relay requests to back-end servers.

Use the front-end and back-end architecture to:

- Maintain a single namespacc to access all front-end servers.
- Eliminate the burden of Secure Sockets Layer (SSL) protocol on back-end servers.
- Provide enhanced security in a firewall or perimeter network environment.

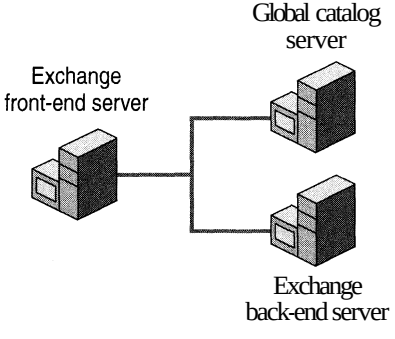


Figure 25.2 Front-end and back-end server architecture

Securing Outlook Web Access on the Internet

The front–end and back-end architecture is useful when enabling users to retrieve their mail from the lnternet. A front–end server can be placed behind or in front of a firewall, or a perimeter network, also called a demilitarized zone (DMZ). A perimeter network provides two layers of filtering, one between the lnternet and the front–end server, and another between the front–end server and the company's network. Different ports on the routers and firewalls must allow access to back-end servers, depending on the location of the front-end server in relation to the firewall.

If you locate servers between the lnternet and the firewall, it is required that more ports be opened, which might compromise security.

By making the firewalls the only servers exposed to the Internet, you can control security with fewer computers and secure the internal network by limiting the number of access points. Figure 25.3 illustrates a deployment using a front–end server inside a firewall.

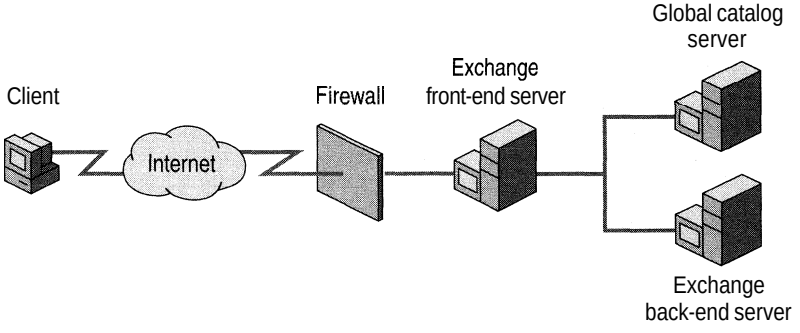


Figure 25.3 Front-end server protected by a firewall

Note that the designation of only one front-end server for multiple back–end servers creates a single point of failure and a possible bottleneck. Monitor this server or deploy more than one front-end server. The front-end server should be configured to use **SSL** to encrypt data and passwords between the client and the front–end server; otherwise passwords and data travel as cleartext.

SSL provides privacy between a Web browser and a Web server. This begins with a handshake phase that negotiates the encryption and establishes a secure session between the client and the server. After this process ends, all data that is sent between the Web browser and the Web server is encrypted.

SSL increases security, but also reduces performance. The installation of a front–end server to handle **SSL** requests removes processor demand on back–end mailbox servers.

In front–end and back–end architecture, servers are optimized to perform single functions. The front-end server processes all the HTTP and **SSL** requests and the back–end server manages data (mailbox stores and public folder stores).

Load Balancing and Fault Tolerance

Network Load Balancing can be configured to balance traffic across multiple front-end servers. Clients access the servers configured with Network Load Balancing by using one Domain Name System (DNS) name, or Internet Protocol (IP) address. Network Load Balancing automatically balances the networking traffic between the servers to create a reliable solution.

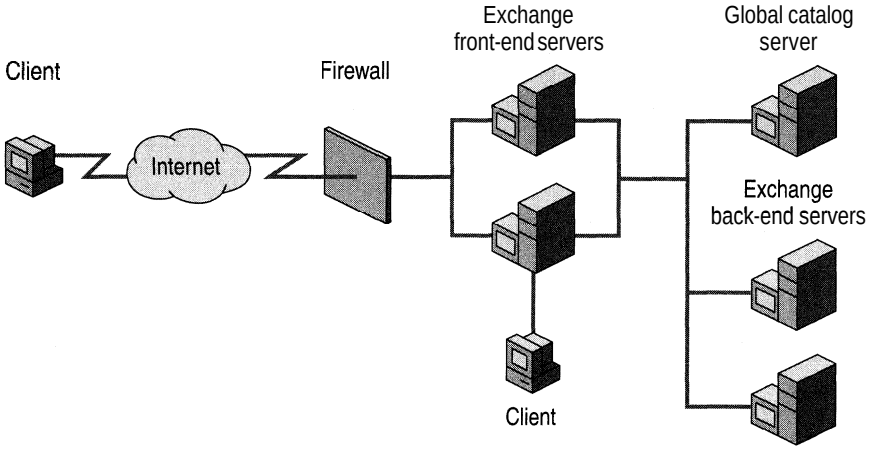


Figure 25.4 Load balancing with multiple front-end servers

The front-end server should be monitored for user traffic by using tools such as Httpmon.exe. If user traffic is too high, Network Load Balancing can be used with multiple front-end servers to distribute the load, thus increasing performance between the client and the front-end server. In this case, the load can be equally distributed across two or more servers. This also provides fault tolerance if a server stops working. Because there are multiple front-end servers, the other can always service the requests.

Note You can download Httpmon.exe from the Microsoft Web site at www.microsoft.com.

Front-End and Back-End Server Architecture

This section is an overview of high-level concepts relating to the configuration of back-end servers and front-end servers. This section also describes how to configure Network Load Balancing for Outlook Web Access.

There are three tasks that need to be configured to allow for front-end and back-end architecture. They include:

- Configuring back-end servers.
- Configuring front-end servers. (This includes configuring Network Load Balancing and installing Exchange 2000 Server.)
- Configuring DNS entries.

For detailed procedures on the configuration of Outlook Web Access, see the Exchange 2000 Server online documentation.

Configuring Back-End Servers

Configure back-end servers to belong to the same domain as the front-end servers. You can configure public folder replication in a way that suits your needs; the back–end server design does not affect public folder replication. Create a virtual server for every front–end server namespace. Virtual servers must match the configuration on the front-end servers.

Make a separate Exchange virtual server for each front-end namespace. For instance, in a front–end namespace consisting of tokyo.microsoft.com and london.microsoft.com, create two virtual servers.

When Exchange runs in a server cluster, Exchange automatically installs protocol virtual server instances on each Exchange virtual server. To run a front-end server with a back–end server cluster, make additional protocol virtual servers on the back–end servers. Ensure that the appropriate static IP is selected for new HTTP virtual servers on Exchange virtual servers.

Configuring Front-End Servers

Before you install Exchange 2000 Server on the front-end servers, install and configure Network Load Balancing.

Configuring Network Load Balancing

Configure Network Load Balancing during your initial network configuration of Windows 2000—or after you configure the server for Outlook Web Access.

Note Network Load Balancing was previously called Windows Load Balancing Service.

To configure Network Load Balancing for Outlook Web Access

1. Open the **Properties** page for **a** network connection object. On the **General** tab, select **Network Load Balancing,** and then click **Properties.**
2. On the **Cluster Parameters** tab, in **Primary IP address,** type the shared IP address for all servers participating in the Network Load Balanced configuration.

   In **Subnet mask,** type the subnet mask for the primary IP address.

   In **Full Internet name,** type the fully qualified domain name (FQDN) for this server.
3. On the **Host Parameters** tab, in **Priority (Unique host ID),** designate the host ID; select a unique ID for each server participating.

   Select **Initial cluster state;** this will make the node join the cluster upon start up.

   In **Dedicated IP address,** type the IP address of the computer; ensure this IP address is unique to this computer and not shared by any other computer on the network.

   In **Subnet mask,** type the subnet mask of the dedicated IP address.
4. On the **Port Rules** tab, in **Affinity,** select **Single.**

   Note This setting is only required when connecting to a front-end server over SSL. It ensures that when a client establishes a session through a front–end server, the client continues to use that server for all requests during that session. If you do this for SSL connections, you reduce the amount of processing and network overhead required to establish new secure sessions.

   In **Load weight** select **Equal.** Click OK.
5. In the **Network Adapter** dialog box, on the **General** tab, select **Internet Protocol (TCP/IP),** and then click **Properties.**
6. On the **Protocol (TCP/IP)** tab, click **Advanced.**
7. On the **Advanced TCP/IP Settings** tab, click **Add,** and then type the shared IP address and the subnet mask of the servers that are configured to participate in Network Load Balancing.

Complete the preceding Network Load Balancing and TCP/IP configuration on all servers that are designated to be front-end servers. Table 25.1 shows which parameters are shared among all nodes, and which parameters are unique.

Table 25.1 Network Load Balancing parameters

Parameter	Shared or Unique
Primary IP address/subnet mask	Shared
Full Internet name	Shared
Priority/Host ID	Unique
Dedicated IP address/subnet mask	Unique
Affinity	Shared (Single)
Load weights	Shared (Equal)

Installing Exchange 2000 Server

After the Network Load Balancing Service is configured on all front-end servers, the server must be configured as a front-end server.

For each front-end server, install Exchange 2000. When the installation is complete, open Exchange System Manager, open **Server Properties,** and then select **This is a front-end server.** Dismount and remove the mailbox stores and public folder stores, and then restart the computer. For information about working with mailbox stores and public folders stores, see the Exchange 2000 Server documentation.

Configuring DNS Entries

For each group of Exchange 2000 front-end servers that use Network Load Balancing, create a DNS entry in the A (host) record. This enables users to access the front–end servers by referring to them with display names rather than by IP address.

Authentication

Outlook Web Access supports several authentication methods. The method that is implemented within your company will vary depending on the client operating system, browser, and security policies.

Basic Authentication

Basic authentication is the most flexible type of authentication and is part of the HTTP specification. It is supported by most browsers and is independent of hardware platform. During the Basic authentication process, the user enters their user name, domain, and password. Specify the domain by typing *domain\username* in the browser's authentication prompt. Basic authentication supports the use of a front-end server. This method is not considered secure because passwords are not encrypted when they are sent to the server. Basic authentication is typically used within a company's intranet or with SSL encryption to protect the password.

Authentication

Integrated Windows Authentication uses the existing credentials that were supplied at log on to access the server. No additional authentication is required. Security is enhanced, because passwords are encrypted when they are sent to the server. Internet Explorer 4.0 or Internet Explorer 5.0 on Windows platforms support Integrated Windows Authentication; other browsers do not. Clients using Windows 2000 authentication use Kerberos. Non-Windows 2000 clients use NTLM protocol for authentication.

> **Note** Integrated Windows Authentication is not supported in the front–end and back-end configuration.

Secure Sockets Layer

To achieve the highest level of security and operability between browser clients and servers, use SSL. Although it provides a secure communications channel between the server and the client, SSL is not an authentication method itself. Basic authentication is commonly used with SSL.

Creating and deleting these SSL sessions places an additional burden on the server. However, if a front-end or back-end configuration is used, front-end servers can manage SSL.

Anonymous Access

Anonymous access allows a user to access a resource without being prompted for authentication. Because of this, it is not necessary to have an account defined for the user.

Front-End and Back-End Authentication

The front–end and back–end architecture presents a more complex security configuration than a standard Outlook Web Access configuration. This section focuses on the two authentication methods that are used within the front–end and back–end architecture: *pass–through* and *dual*.

Pass-Through Authentication

With pass-through authentication, the front-end server supports anonymous access. This allows clients to send the authentication request through the front-end server to the back-end server. The back-end server prompts the client for credentials and receives the response.

Dual Authentication

Dual authentication requires that both the front-end and back-end servers authenticate the user. However, only the front-end servers prompt the client for authentication credentials. After the front-end server receives the client's credentials, the back-end server receives the same information and does not reissue the prompt for the client's credentials. Dual authentication provides the front-end server the ability to determine the public folder server associated with a user's mailbox without querying a back-end server.

Performance Monitoring

You can use System Monitor to determine the scalability of the Outlook Web Access component of Exchange 2000. If you understand the performance criteria that have a direct relationship with system performance, you can monitor performance to:

- Maximize system performance.
- Improve a system with unacceptable performance.
- Measure the overall system load.

To understand how to increase the performance of the system, identify the factors that limit performance. After you do so, you can take corrective action to increase system performance.

The architecture of the Outlook Web Access component of Exchange 2000 depends on the Web Storage System. The limitations of the system exist in the system hardware rather than a specific application or protocol. Therefore, CPU use, memory, disk subsystem, and network performance all limit performance.

Monitor the objects listed in Table 25.2 to clarify your system's limitations.

Table 25.2 Objects and their identifiers

Object	Parameter
Processor	% Processor Time
System	Context Switches/Sec
Process	% Process Time/Store % Process Time/inetinfo % Process Time/lsass % Process time/mad
Physical disk	Disk Reads/sec Disk Writes/sec Current Disk Queue Length
Memory	Available Bytes Page Reads/sec Page Writes/sec Page Faults/sec

Because the performance limitations of Outlook Web Access relate to the hardware, as system hardware capacity increases, so does the performance and scalability of Outlook Web Access.

Usage patterns may also be monitored with the Exchange 2000 Server HTTP Extension object in System Monitor.

Applications

Windows 2000 includes a number of applications that allow system administrators to monitor performance. These applications include:

- Performance Tool
- Event Viewer
- Task Manager
- Network Monitor
- Httpmon.exe

The Windows 2000 Performance Tool is composed of two parts: System Monitor and Performance Logs and Alerts. The System Monitor component allows administrators to collect and view real-time data about memory use, disk subsystem, processor use, network performance, or other activities. This information can be evaluated in graph, histogram, or report formats. The Performance Logs and Alert component allows you to configure logs to record performance data and set system alerts to notify you when a specified counter's value is above or below a defined threshold. This portion of the performance tool is often used to notify administrators of system problems before users experience the problem.

Httpmon.exe (Httpmon) is a utility included with the *Microsoft Windows* 2000 *Server Resource Kit Internet Information Services 5.0 Resource Guide.* It emulates how users attempt to connect to Web sites. Httpmon can run on several servers simultaneously to test connectivity and display the connectivity statistics of the sites in your test site list. Httpmon is a multi-threaded process that operates according to the parameters and values set in its associated Httpmon.ini file. These parameters specify the Web sites to be tested, in addition to the number of times and frequency to retry a failed connection before moving on to the next specified site.

Although Windows 2000 Server includes several applications to monitor system and application performance, there are also several third–party applications that specialize in event monitoring. Knowing the type of information you need to monitor and how you want that information reported is the evaluation criteria for choosing a third–party application.

APPENDIX A

Client Network Traffic Analysis

Christophe Leroux, Consultant, Microsoft
Christophe Besançon, Consultant, Microsoft

Deploying Microsoft Exchange 2000 Server in a company requires a thorough understanding of a company's structure, network topologies, usage patterns, and so on. A particularly difficult task is estimating the traffic generated between multiple Microsoft Exchange 2000 servers and messaging clients.

This chapter provides a detailed look at how Microsoft Exchange 2000 consumes bandwidth under different configurations and analyzes the traffic between clients and Microsoft Exchange 2000 Server.

System professionals can use the statistics and information to better understand network bandwidth loading, including how to set up segments, assign users, and extrapolate the results for various configurations.

The traffic analysis in this chapter is based on the results of 700 tests conducted with the following messaging clients:

- Microsoft Outlook 2000
- Microsoft Outlook 97
- Microsoft Outlook Express 5.0
- Microsoft Outlook Web Access
- Microsoft Exchange Instant Messaging 2.0
- Netscape Communicator 4.7

The data from these tests appears in Microsoft Excel workbooks on the *Microsoft Exchange* 2000 *Sewer Resource Kit* companion *CD*. In some cases, the data captures do not represent all information required by the clients. Network bandwidth is not consumed for cached graphics, session lifetimes, cached resolved names, and message formats. However, message content and form significantly impacts network traffic, especially an HTML message. These tests allow you to compare clients and mail message format Rich Text Format (RTF), HTML, Plain Text. You can also use the data to estimate the network traffic generated by many user profiles when their messaging habits are well known.

> **Note** Microsoft Outlook 98 was not tested because its performance is similar to the performance of Outlook 2000. For more information about testing Outlook 98, see "Directory Access" later in this chapter.

In This Appendix

Test Lab Configuration

Log On and Log Off

Directory Access

Mail Items

Calendaring, Contacts, and Tasks

Public Folders

Outlook 2000 with Terminal Services

Web Storage System

Instant Messaging

Client Traffic Measurement Conclusions

DSProxy

Test Lab Configuration

The test lab was built with four computers running Windows 2000 Server. One server ran Exchange 2000 Server RC2 (build 4386) and the others ran Windows 2000 services like domain controller, global catalog, and Domain Name System (DNS).

The three client computers had the following products installed:

- Outlook 2000 (version 9.0.0.2711)
- Outlook 97 (version 8.04.5619)
- Outlook Web Access (with Microsoft Internet Explorer 5.0)
- Outlook Express 5.0 (version 5.00.2919.67.00)
- Netscape Messenger 4.7
- Exchange Instant Messaging 2.0 (version 2.0.1002)
- Terminal Services client (32-bit version)

Figure A.1 illustrates the test lab configuration.

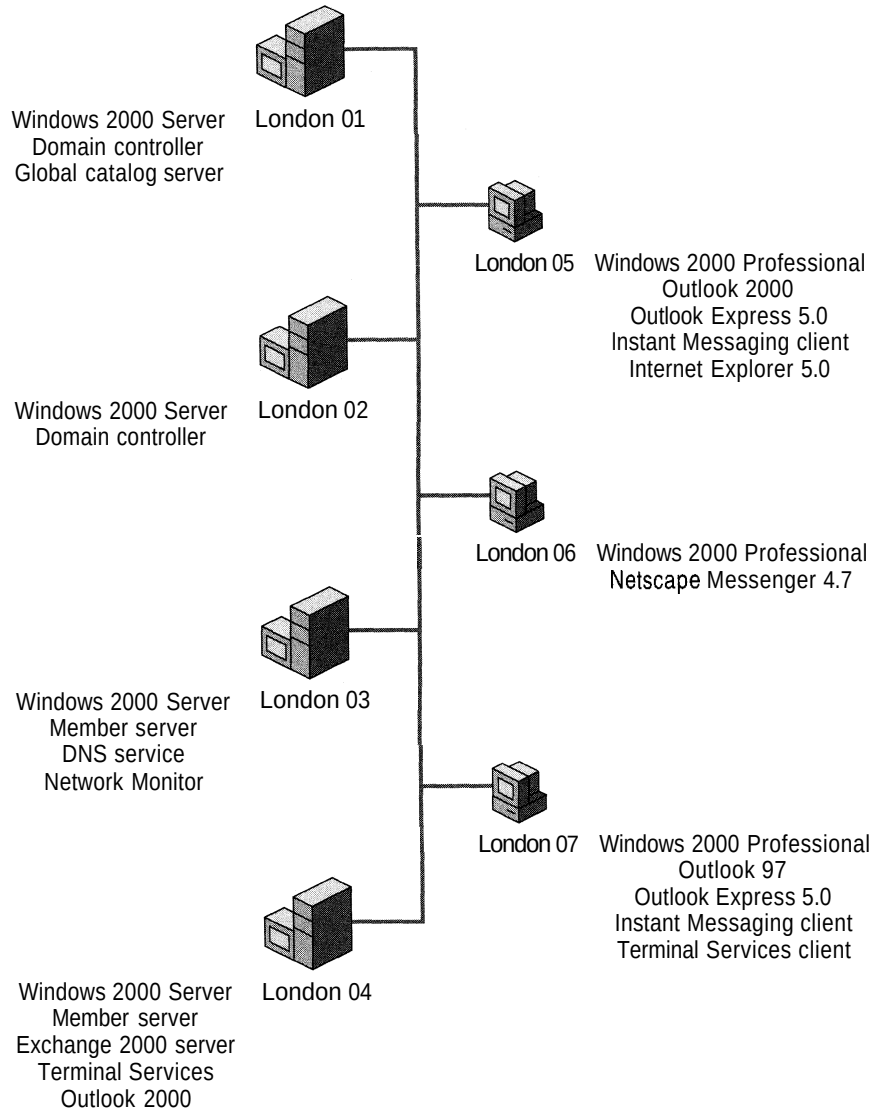


Figure A.1 Test lab network topology

LAN Design

The DNS root domain is called *microsoft.com.* The server, LONDON–03, manages the DNS primary zone and allows dynamic updates. Every server has a static Internet Protocol (IP) address and refers to the DNS to resolve names.

To avoid disturbing the tests and experiencing caching phenomena, the media access control (MAC) and IP addresses are preloaded on the client computers. This prevents Address Resolution Protocol (ARP) or DNS queries from being included in the test results.

The network is an Ethernet 100–megabyte (MB) dedicated network.

Table A.1 Sewer configurations

Server	IP Address	IP Subnet Mask	Role
LONDON-01	10.0.0.1	255.0.0.0	Domain controller Global catalog
LONDON-02	10.0.0.2	255.0.0.0	Domain controller
LONDON-03	10.0.0.3	255.0.0.0	Member server DNS
LONDON-04	10.0.0.4	255.0.0.0	Member server Exchange 2000 Server
LONDON-05	10.0.0.5	255.0.0.0	Member workstation
LONDON-06	10.0.0.6	255.0.0.0	Member workstation
LONDON-07	10.0.0.7	255.0.0.0	Member workstation
LONDON-08	10.0.0.8	255.0.0.0	Member server

Server Characteristics

The following table details the processor speed, memory, and software that runs on each server.

Table A.2 Server characteristics

Server	CPU	Memory	Operating System	Other Software
LONDON-01	Intel Pentium II 450 MHz	128 MB	Windows 2000 Advanced Server	
LONDON-02	Intel 486 66 MHz	64 MB	Windows 2000 Advanced Server	
LONDON-03	Intel Pentium III 500 MHz	128 MB	Windows 2000 Advanced Server	
LONDON-04	Intel Pentium III 500 MHz	128 MB	Windows 2000 Advanced Server	Exchange 2000 Server
LONDON-08	Intel Pentium III 500 MHz	128 MB	Windows 2000 Advanced Server	Terminal Services

Client Characteristics

The following table details the processor speed, memory, and software that runs on each client computer.

Table A.3 Client characteristics

Client	CPU	Memory	Operating System	Other Software
LONDON-05	Intel Pentium 166 MHz	144 MB	Windows 2000 Professional	Outlook 2000 Outlook Express 5.0 Instant Messaging 2.0
LONDON-06	Intel Pentium 166 MHz	144 MB	Windows 2000 Professional	Netscape Messenger 4.7
LONDON-07	Intel Pentium III 500 MHz	128 MB	Windows 2000 Professional	Outlook 97 Outlook Express 5 Instant Messaging 2 Terminal Services client

Exchange 2000 Configuration

These tests capture network traffic between a client and a server. Use only one Exchange 2000 server for this set of tests.

Install the server in the default routing group as a custom setup, because Instant Messaging Service is also installed.

The following screen shot summarizes the installed components on this server.

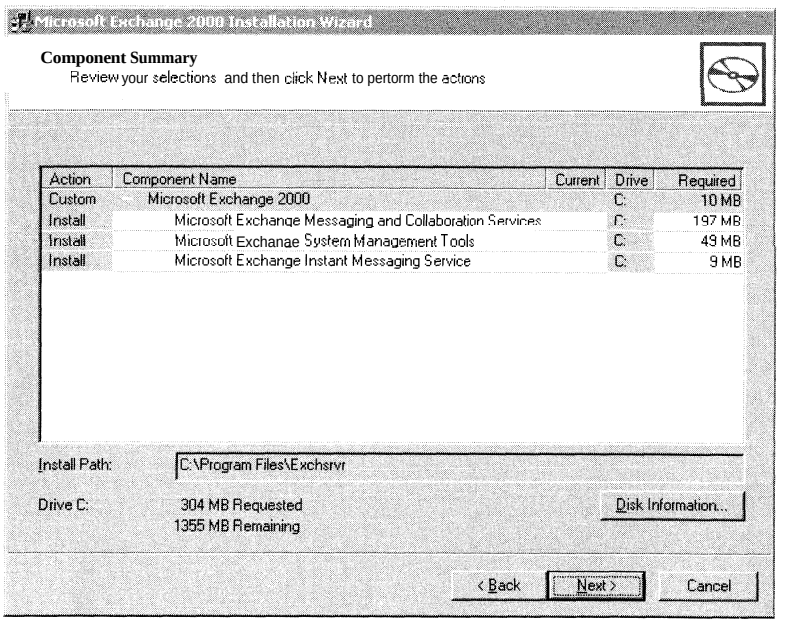


Figure A.2 Exchange 2000 component summary

The messaging organization is LitWare (a fictional company), and the server is in the default administration group, First Administrative Group. The mailbox and public folder stores appear in the only storage group, First Storage Group. No server options, mailbox store options, or public folder store options have been changed; all of the settings are defaults.

Create users in the default Windows 2000 organization unit. The test users have a mailbox, an e-mail address, and no password. Because Exchange Instant Messaging demands a password, only Instant Messaging users have a password.

Terminal Services Configuration

The LONDON-08 server is running Terminal Services and Terminal Services Licensing. This enables the multi-user environment to access the server; thus, a user can open a session remotely and run any installed server application. The server evaluates the network traffic when using Outlook 2000 in the Terminal Services environment.

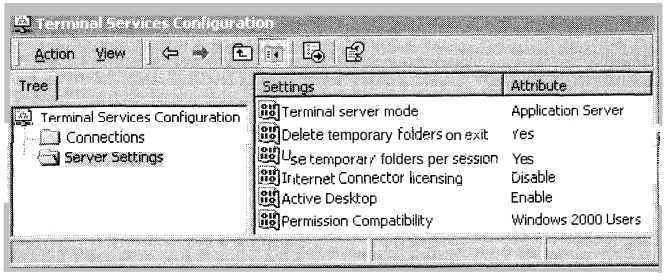


Figure A.3 Terminal Services configuration

Client Configuration

For the tests, use the default client configuration defined by Setup. However, you can alter some settings to reduce disruption in captures or to make the client comparisons easier.

Outlook 2000 Client Configuration

The Outlook 2000 default MAPI test profile included only two services: Microsoft Exchange 2000 Server and Microsoft Outlook Address Book.

The following client settings were cleared before testing:

- **Empty the Deleted Items folder upon exiting**
- **Save copies of messages in Sent Items folder**
- **AutoSave**
- **Preview Pane**

The folder list is enabled only for the tests on public folders.

Outlook 97 Client Configuration

The MAPI test profile included only three services: Microsoft Exchange Server, Outlook Address Book, and Personal Address Book.

The following client settings were cleared:

- **Empty the Deleted Items folder upon exiting**
- **Save copies of messages in Sent Items folder**

The folder list is enabled for the tests on public folders.

Outlook Express Client Configuration

The following Outlook Express settings were cleared:

- Check for new messages every *x* minutes
- Automatically log on to MSN Messenger Service
- Save copy of sent messages in Sent Items folder
- Purge deleted messages when leaving IMAP folders

The settings Send and receive messages at startup and When starting, go directly to My **Inbox** folder will be described in each test.

The HTML format is embedded in Multipurpose Internet Mail Extensions (MIME) format and encoded with Quoted Printable.

The Plain Text format is encoded as MIME with Quoted Printable.

The Post Office Protocol version 3 (POP3) and Internet Message Access Protocol version 4 (IMAP4) accounts are created alternatively, so that the client never supports more than one mailbox.

POP3 Account Settings

- The setting Include this account when receiving mail or synchronizing is selected.
- The server with POP3 and Simple Mail Transfer Protocol (SMTP) is LONDON–04.
- The setting Log on using Secure Password Authentication is cleared.
- The setting Leave a copy of message on server is cleared.

IMAP4 Account Settings

- The setting Include this account when receiving mail or synchronizing is selected.
- The server with IMAP4 and SMTP is LONDON–04.
- The setting Log on using Secure Password Authentication is cleared.
- The setting Break apart messages larger than is cleared.

Lightweight Directory Access Protocol Account Settings

- The Lightweight Directory Access Protocol (LDAP) server is LONDON–01.
- The setting This server requires me to log on is selected.
- The account name is cn=alias, cn=microsoft.
  ("cn=alias,cn=microsoft.com", "cn=alias,dc=microsoft.com" also works).

- The setting Log on using Secure Password Authentication is selected; otherwise, the log on will fail.
- The box Search base contains: dc=Microsoft, dc=com.
- The setting Use simple search filter is cleared.

Network News Transfer Protocol Account Settings

The Network News Transfer Protocol (NNTP) server is LONDON–04.

- The setting This server requires me to log on is cleared. The default virtual-NNTP server allows anonymous access.

> **Note** Create the NNTP account only with the POP3 account. IMAP4 allows native connection to public folders; POP3 does not support native connection.

Netscape Messenger Client Configuration

The following options under Mail and Newsgroups are disabled:

- Addressing, Address Books, Directory Server.
- All options in Copies and folders.
- All options in Formatting.

The following options are enabled:

- In Messages, Send messages that use 8-bit characters, using the Quoted Printable MIME encoding.
- In Formatting, Send the message in HTML anyway.

Create the POP3 and IMAP4 accounts alternatively so the client never supports more than one mailbox. Netscape Messenger does not allow a POP3 account while there is an IMAP4 account; it does, however, allow multiple IMAP4 accounts.

POP3 Account Settings

- The server with POP3 and SMTP is LONDON-04.
- The option Check for mail every x minutes is disabled.
- The option Leave a copy of message on server is enabled by default.

IMAP4 Account Settings

- The server with IMAP4 and SMTP is LONDON-04.
- The option Check for mail every x minutes is disabled.
- The option Clean up **Inbox** on exit is disabled.

NNTP Account Settings

- The NNTP server is LONDON-04.
- The option **Only ask me for my user name and password when necessary** is enabled. The logon is anonymous.

> **Note** Netscape Messenger creates an NNTP account automatically. Because accessing public folders is possible with IMAP4, use the NNTP account with the POP3 account.

Outlook Web Access Client Configuration

In Outlook Web Access tests, we chose to use Microsoft Internet Explorer 5 to get the benefits of the XML support.

The amount of network traffic varies, depending on whether or not graphics are enabled on the browser. When graphics are enabled and present on HTML pages, Internet Explorer searches for graphics in its cache folder before downloading. If a graphic is found in the cache, there is less network traffic. The data for the tests reflect normal Outlook Web Access behavior. Therefore, graphics are loaded from the cache unless indicated otherwise for a particular test.

Terminal Services Client Configuration

The Terminal Services 32-bit client is installed on a workstation running Windows 2000 Professional.

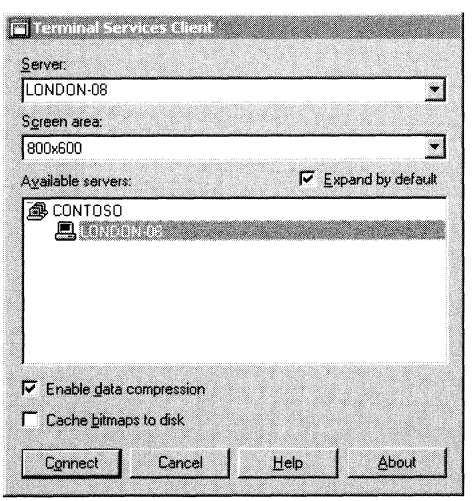


Figure A.4 Terminal Services Client

The default Terminal Services client parameters are illustrated in Figure A.4. The screen area is set to 800 x 600 pixels and the **Enable data compression** check box is selected.

Measurement Methodology

The network traffic captures were done with Microsoft Network Monitor 2.0 version 5.00.943, found in Microsoft System Management Server 2.0 Service Pack 2.

Figure A.5 shows the test capture filter.

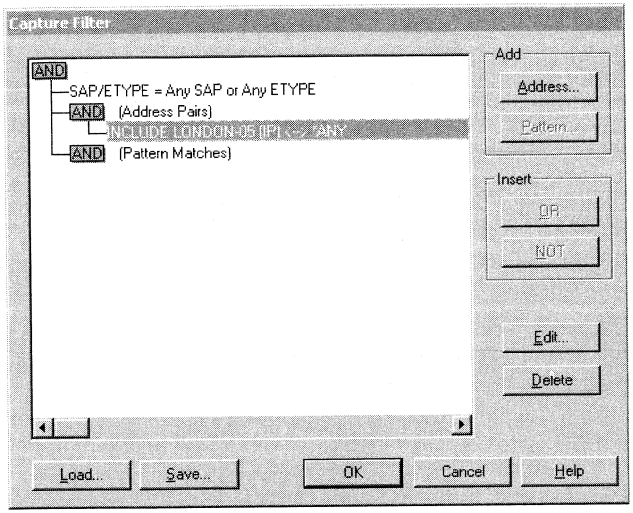


Figure A.5 Network Monitor Capture Filter

A static network address resolution was established to prevent unexpected network broadcasts for MAC addresses or requests from DNS to get IP addresses from name, both of which cause timeouts.

A batch file containing the command Arp.exe –s *IP address* Mac *address* was created for each computer in the lab network. This batch file ran on each computer with the following command:

```
ARP.EXE -s 10.0.0.1 01-02-03-qb-cd-ef-99
```

A host file that contained all computer names and their IP address was also created on each computer.

Each data capture was performed according to the following steps:

1. Start capture.
2. Perform the action to measure.
3. Wait until there is no more traffic.
4. Stop capture.
5. Save the capture.
6. Report results.

Each capture was performed two or three times for results comparison and consistency checks.

The test messages had the following characteristics.

For all the tests, except the public folders tests:

- Messages of 1 kilobyte (KB), 2 KB, and 4 KB were created, their sizes determined by the size appearing on Outlook 2000 when the messages were sent using RTF. These messages were saved as .msg files and their size on the hard drive was greater than when they were stored in the mailbox. On the hard drive, sizes were 5 KB, 10 KB, and 20 KB, respectively. These messages contained a few characters per row and many carriage returns (message content impacts network traffic when the messages are converted into HTML format). The font in each message was Arial. The font style was regular, and the font size was 10 points.
- Messages with attachments: had Microsoft Word or Microsoft PowerPoint files of 10 KB, 50 KB, 100 KB, 500 KB, and 1,000 KB attached to the 1–KB message described earlier.

Public folders tests and NNTP access:

- Messages for public folders tests contained 1024, 2048, or 4048 characters, creating messages of 1 KB, 2 KB, and 4 KB respectively. These non-RTF messages were saved as .txt files on a hard disk drive. The text contained a portion of the Microsoft Excel 97 Readme.txt and it contained no more than 75 characters per line.
- Messages with attachments: Microsoft Word or Microsoft PowerPoint files of 10 KB, 50 KB, 100 KB, 500 KB, and 1,000 KB, with the text of the 1–KB message (1024 characters) described earlier.

Log On and Log Off

These tests measure the network traffic generated by initial client connection and validation on the Microsoft Exchange 2000 Server when users access their mailboxes. Because each client is unique, the tests are partitioned to the following four groups.

- MAPI clients: Microsoft Outlook 2000, Outlook 97
- Microsoft Outlook Express: POP3 and IMAP4 modes
- Netscape Messenger: POP3 and IMAP4 modes
- Web client: Microsoft Outlook Web Access

Tests Performed

Log On

This test captured the traffic generated by initial client connection and validation and that of Microsoft Exchange Server to the global catalog server and domain controller. For MAPI clients (Outlook 2000 and Outlook 97), the traffic was captured from when Outlook was launched until the traffic dissipated.

Log Off

This test captured the traffic resulting from a disconnection from the server.

MAPI Clients: Microsoft Outlook 2000, Outlook 97

For MAPI clients, the generated traffic depends on various parameters:

- **Whether Outlook is running for the first time on the computer** If so, one or more welcome messages arrive in the mailbox.
- **Whether the mailbox is new** Mailboxes created on a server are not formatted for Outlook. When you connect with Outlook to a new mailbox, extra traffic is generated to initialize the mailbox on the server-that is, to create system folders (Calendar, Contacts, Drafts, Journal, Notes, Tasks) and system views. Only Outlook 2000 creates a Drafts folder.
- **Whether the MAPI profile is new** When a MAPI profile is new, extra traffic is generated to initialize it with settings stored in the registry for domain names of the resolved name, the mailbox servers, and public folder servers.
- **Whether the preview pane view is on** Only Outlook 2000 has this feature. The preview pane is activated by default. Although the preview pane view can be added to Outlook 97, the tests described in this appendix did not include Outlook 97 with the preview pane view.
- **Whether default folder contains items** The item count in the default folder has an impact on traffic. Captures were done with inboxes that contained 50 and then 100 unread items.

Test Details

The following tests were performed.

LOGON1: First startup, new mailbox, new profile

1. Create a new mailbox on the Exchange server.
2. Install Outlook on the client computer.
3. Create a profile with the server name and user name.
4. Start Outlook.

To capture the traffic, the user name was not resolved.

LOGON2: First startup, existing mailbox, new profile

1. Install Outlook on the client computer.
2. Create a profile with the server name and user name.
3. Start Outlook.

In order to capture the traffic, the name was not resolved.

A MAPI client used the mailbox, but it was empty.

LOGON3: New mailbox, new profile

1. Create a new mailbox on the Exchange server.
2. Create a profile with the server name and user name.
3. Start Outlook.

To capture the traffic, the name was not resolved.

LOGON4: Existing mailbox, new profile

In this test, a MAPI client had used the mailbox, but the mailbox was empty.

1. Create a profile with the server name and user name.
2. Start Outlook.

To capture the traffic, the name was not resolved.

LOGON5: Existing mailbox, existing profile, no message

This is a common situation. A profile and mailbox exist, and a MAPI client has used the mailbox already, but the mailbox is now empty. In this test, the preview pane was disabled.

- Start Outlook.

LOGON6: Existing mailbox, existing profile, 50 messages

In this test, the preview pane was disabled so that the size of the messages did not affect the capture, and so that results between Outlook 2000 and Outlook 97 could be compared.

1. Start Outlook.
2. Send 50 messages with 1024 characters as content.
3. Exit and log off Outlook.
4. Start Outlook.

Inbox contains 50 unread messages.

LOGON7: Existing mailbox, existing profile, 100 messages

In this test, the preview pane was disabled so that the size of the messages did not affect the capture, and so that results between Outlook 2000 and Outlook 97 could be compared.

1. Start Outlook.
2. Send 100 messages with 1024 characters as content.
3. Exit and log off Outlook.
4. Start Outlook.

Inbox contains 100 unread messages.

LOGON8: First startup, existing mailbox, new profile, no preview pane

This test was done with Outlook 2000 to measure the preview pane's affect on the network. This test was done on an existing mailbox.

1. Run Outlook 2000.
2. In Inbox, disable preview pane.
3. Uninstall Outlook 2000.
4. Install Outlook 2000.
5. Create a profile with the server name and user name.
6. Run Outlook.

To capture the traffic, the name was not resolved.

LOGOFF: Logoff

1. Close all Outlook windows except the main one.
2. Close Outlook by clicking **File,** and then click **Exit and Log Off**

Log On and Log Off Results

For the tabulated results of the log-on and log-off tests for each mail client, see LogOnLogOff.xls on the *Exchange* 2000 ***Server Resource Kit*** companion CD.

Outlook 2000 and Outlook 97 Measurements Analysis

The measurements captured when running Outlook 2000 and Outlook 97 for the first time are not comparable. The number and size of welcome messages and the preview pane (only present in Outlook 2000) explains the difference.

The creation of a new profile or new mailboxes generates more traffic. A new profile generates 30 KB more traffic. The initialization of a new mailbox generates 3 KB more traffic. In daily use (where there is an existing mailbox and profile), Outlook 2000 generates as much traffic (in bytes) as Outlook 97.

As more fields appear, such as previews of unread messages, network traffic increases. The size of the message does not affect the traffic.

A close analysis of the data shows that Outlook 2000 contacts the closest global catalog server, whereas Outlook 97 does not. Outlook 97 contacts the Exchange server, which then sends the directory look-up requests to the global catalog server. The global catalog server is contacted only during the logon process and during address book-related queries.

Both versions of the Outlook clients generate equal network traffic when closing. Outlook 2000 and Outlook 97 consume the same bandwidth during logon and logoff.

Microsoft Outlook Express: POP3 and IMAP4 Modes

For Outlook Express, the traffic generated depends on various parameters:

- **Whether Outlook Express is configured to automatically check for new messages at logon** Outlook Express allows users to select whether or not to connect to the server when starting.
- **What quantity, size, and structure of the messages is to be downloaded** The quantity, size, and structure (the presence or absence of an attachment) of the messages affect the traffic.
- **Whether Outlook Today is displayed by default**

Test Details

The following network traffic tests were performed on Outlook Express.

LOGON1: "Send and receive message at startup" option disabled

1. Create a POP3 or IMAP4 account in an empty mailbox.
2. Clear the Send and receive message at startup setting.
3. Run Outlook Express.

LOGON2: Inbox, No message, "Send and receive message at **startup"** option enabled

1. Create a POP3 or IMAP4 account empty mailbox.
2. Select the option Send and receive messages at startup.
3. Select the option When starting, go directly to my **'Inbox'** folder.
4. Run Outlook Express.

LOGON3, LOGON4, and **LOGON5: Inbox,** 1-KB, 5-KB, and 10-KB messages, "Send and receive message at **startup"** option enabled

1. Create a POP3 or IMAP4 account empty mailbox.
2. Select the option Send and receive message at startup.
3. Select the option When starting, go directly to my **Inbox** folder.
4. Send 1–KB, 5–KB, and 10–KB messages.
5. Run Outlook Express.

LOGON6: Outlook Today, No message, "Send and receive message at startup" option enabled

1. Create a POP3 or IMAP4 account empty mailbox.
2. Select the option Send and receive message at startup.
3. Clear the option When starting, go directly to my **'Inbox'** folder.
4. Run Outlook Express.

LOGON7, LOGON8, LOGON9: Outlook Today, **1-KB,** 5-KB, 10-KB messages, "Send and receive message at startup" option enabled

1. Create a POP3 or IMAP4 account empty mailbox.
2. Select the option Send and receive message at startup.
3. Clear the option When starting, go directly to my **'Inbox'** folder.
4. Send 1–KB, 5–KB, and 10–KB messages.
5. Run Outlook Express.

LOGOFF

1. Close all Outlook windows, except the main window.
2. Close Outlook by clicking **File,** then click **Exit and Log Off.**

Outlook Express IMAP4 and POP3 Results

For the tabulated results of the log-on and log-off tests for each mail client, see LogOnLogOff.xls on the *Exchange* 2000 *Server Resource Kit* companion CD.

Outlook Express Measurements Analysis

When comparing an Outlook Express IMAP4 or POP3 client, keep in mind that IMAP4 and POP3 are two transport protocols with different features. IMAP4 is more powerful and advanced than POP3.

With an IMAP4 account, launching Outlook Express always generates 2 KB of traffic to the server because it opens a connection with the Exchange 2000 server to receive incoming message notification. With a POP3 account, Outlook Express generates traffic only if a Send/Receive action is performed.

With IMAP4, displaying Outlook Today consumes less bandwidth than displaying the Inbox. When Outlook Today is displayed, the client does not download headers (sender, subject, received date). The number of messages and the contents affect the traffic generated when displaying the Inbox folder at start up.

With POP3, there is no difference in network traffic whether Outlook Today is displayed or not. POP3 always stores messages locally; Outlook Today reads the messages stored on the hard drive. The only traffic in this situation is generated during connection, and when messages are download.

IMAP4 uses 200 bytes to disconnect. In POP3, connection, download, and disconnection are included in the Send/Receive process.

Overall, POP3 produces less network traffic than IMAP4, but does not offer the same amount of functionality. POP3 can only download messages from one folder. IMAP4 can synchronize many folders stored on a server. Outlook Express (with an IMAP4 account) stores, by default, three folders on the server: Inbox, Drafts, and Sent Items.

Netscape Messenger: POP3 and IMAP4 Modes

When a POP3 account is defined, Netscape Messenger does not offer an option to download new messages at startup. Therefore, no traffic occurs. No capture occurred in POP3 mode.

When an IMAP4 account is defined, Netscape Messenger automatically checks for new messages at startup.

For Netscape Messenger, the generated network traffic depends on the quantity, size, and structure of the messages to be downloaded.

Test Details

The following tests were performed.

LOGON2: Inbox, no new message

1. Create an IMAP4 account in an empty mailbox.
2. Run Netscape Messenger.

LOGON3, LOGON4, LOGON5: Inbox, 1-KB, 5-KB, and 10-KB messages

1. Send 1-KB. 5-KB, and 10-KB messages.
2. Create an IMAP4 account in an empty mailbox for each test.
3. Run Netscape Messenger for each test.

LOGOFF

1. Close all Netscape windows except the main one.
2. Close Netscape Messenger by pressing ALT+F4.

Netscape Messenger IMAP4 Results

For the tabulated results of the log-on and log-off tests for each mail client, see LogOnLogOff.xls on the Exchange 2000 ***Server Resource Kit*** companion CD.

Netscape Messenger IMAP4 Measurements Analysis

By default, Netscape Messenger stores only the Inbox on the server and requires that the user request new messages. Outlook Express opens a connection at log on to receive notification when new messages arrive. Thus Outlook Express uses more bandwidth to log on.

Outlook Web Access

For Outlook Web Access, the traffic generated depends on the quantity, size, and structure of the messages to be downloaded.

Test Details

The following tests were performed.

LOGON2: Inbox, no new message

1. Create a mailbox.
2. Run Outlook Web Access (http://LONDON-04/exchange).

LOGON3, LOGON4, LOGONS: Inbox, 1-KB, 5-KB, and 10-KB new messages

1. Send 1–KB, 5–KB, and 10–KB messages.
2. Run Outlook Web Access (http://LONDON-04/exchange).

LOGOFF

1. Run Outlook Web Access (http://LONDON-04/exchange).
2. Close all Outlook Web Access windows.
3. Close Explorer.

Outlook Web Access Results

For the tabulated results of the log-on and log-off tests for each mail client, see LogOnLogOff.xls on the *Exchange* 2000 *Server Resource Kit* companion CD.

Outlook Web Access Measurements Analysis

A quick analysis shows that Outlook Web Access generates more traffic than other clients. Deeper analysis shows that part of the increased traffic comes from sending the Web interface pages. The number of additional messages in the Inbox creates a small amount of traffic (about 600 bytes per message).

Disconnecting from Exchange generates only 54 bytes of traffic, the least amount of traffic among the tested messaging clients.

Directory Access

Exchange 2000 introduces a new way to resolve address names. With the Active Directory directory service, Exchange 2000 refers address query resolution to the closest global catalog server.

The MAPI clients supporting this direct access method to the global catalog server are:

- Microsoft Outlook 2000 (any version).
- Microsoft Outlook 98 version 8.5.6204.0 and later. This includes the Outlook 98 Archive Patch, available for download from the Microsoft Web site at http://www.microsoft.com.

With earlier MAPI clients, Exchange 2000 forwards queries to the global catalog server. Exchange 2000 server communicates with the global catalog server on behalf the MAPI client.

Earlier MAPI clients are:

- Microsoft Outlook 98 versions before 8.5.6204.0.
- Microsoft Outlook 97 (any version).
- Microsoft Exchange client (any version).

Note Most companies that use Outlook 98 apply the Outlook 98 Archive Patch. With this update, Outlook 98 performs like Outlook 2000. Therefore, measurements are included for Outlook 97 97 8.04.5619 (from SR2b) and Outlook 2000.

Tests Performed

These tests measured the network traffic generated when users check addresses or names and other kinds of access information to the address book. Because the clients differ, the tests are adapted to the four groups:

- MAPI clients: Microsoft Outlook 2000, Outlook 97
- Microsoft Outlook Express: POP3 and IMAP4 mode
- Netscape Messenger: POP3 and IMAP4 mode
- Web Client: Microsoft Outlook Web Access

The following concepts and their corresponding acronyms are referenced in this section. The tests in this section are identified by the acronym that follows each concept below.

Address Resolution (AR) When users check names or use automatic resolution on a recipient.

Ambiguous Name Resolution (ANR) When a name is ambiguous, the user must choose from a list of names. A Check Names action is performed before any ANR tests to prevent excess traffic during initial access to the address book in the Outlook session.

Address Lookup (AL) The address book dialog box appears when users click **To** (or **Cc, Bcc**) in the Outlook client. The **To** button is clicked when addressing a message to capture traffic resulting from users running MAPI queries of addresses in the global address list (GAL).

Address Book View Lookup (ABVL) When users scroll through the GAL in the address book.

Address Details (AD) When users ask for a GAL entry, or a name's properties.

MAPI Clients: Outlook 2000, Outlook 97

For MAPI clients, the traffic generated when accessing the directory depends on various parameters:

- Whether or not the mailbox is new When a new Outlook client has connected to a mailbox, does the first resolution generate extra traffic?
- Whether or not a profile is new When a profile has not yet been used, does the first resolution generate extra traffic?
- How address resolution is accomplished Resolution can occur automatically or by clicking **Check Names.**
- Whether an address is ambiguous When the name is ambiguous, many results return. How does this affect network traffic? A non-ambiguous name or alias returns only one result.
- Whether the name has been resolved prior to this instance It is possible to track a cache optimization.

Test Details

The following network traffic tests were performed.

AR1: Check Names, first time in the profile (first Outlook session)

1. Create a new profile.
2. Start Outlook.
3. Create a new message.
4. Enter a non-ambiguous alias.
5. Click **Check Names.**

AR2: Check Names, existing profile, first Outlook session

1. Start Outlook.
2. Create a new message.
3. Enter a non-ambiguous alias.
4. Click **Check Names.**

AR3: Check Names, second Outlook session

This test preparation occurred after AR2.

1. Create a new message.
2. Enter a non-ambiguous alias distinct from the alias entered in test AR2.
3. Click **Check Names.**

AR4: Check Names, name previously checked

This test preparation was performed after AR3.

1. Create a new message.
2. Enter the non-ambiguous alias entered in test AR3.
3. Click **Check Names.**

AR5: Automatic name checking, first time in the profile (first Outlook session)

1. Create a new profile.
2. Start Outlook.
3. Create a new message.
4. Enter a non-ambiguous alias.
5. Click in **Subject.**

AR6: Automatic name checking, existing profile, first Outlook session

1. Start Outlook.
2. Create a new message.
3. Enter a non-ambiguous alias.
4. Click in **Subject.**

AR7: Automatic name checking, second Outlook session

This test preparation was performed after AR6.

1. Create a new message.
2. Enter a non-ambiguous alias distinct from the alias entered in test AR6.
3. Click in **Subject.**

AR8: Automatic name checking, name previously checked

This test preparation was performed after AR7.

1. Create a new message.
2. Enter the non-ambiguous alias entered in test AR7.
3. Click in **Subject.**

ANR1: Check Names, two results generated

1. Start Outlook.
2. Create a new message.
3. Enter an ambiguous alias that will generate two results.
4. Click **Check Names.**

ANR2: Check Names, five names returned

1. Start Outlook.
2. Create a new message.
3. Enter an ambiguous alias that will generate five results.
4. **Click Check Names.**

ANR3: Check Names, one of two names previously chosen

This test preparation was performed after ANRI.

1. Start Outlook.
2. Create a new message.
3. Enter the ambiguous alias entered in test ANR1.
4. Click **Check Names.**

ANR4: Automatic name checking, two names returned

1. Start Outlook.
2. Create a new message.
3. Enter an ambiguous alias with two results.
4. Click in **Subject.**

ANR5: Automatic name checking, five names returned

1. Start Outlook.
2. Create a new message.
3. Enter an ambiguous alias that will generate five results.
4. Click in the **Subject box.**

ANR6: Automatic resolution, one of two names previously chosen

This test preparation was performed after ANR4.

1. Start Outlook.
2. Create a new message.
3. Enter the ambiguous alias entered in test ANR4.
4. Click in **Subject.**

AL1: To button, first time in the profile (first Outlook session)

1. Create a new profile.
2. Start Outlook.
3. Create a new message.
4. Click **To.**

AL2: To button, first Outlook session

This test preparation was performed after AL1.

1. Start Outlook.
2. Create a new message.
3. Click **To.**

AL3: To button, second Outlook session

This test preparation was performed after AL2.

1. Create a new message.
2. Click **To.**

ABVL1: Scroll down one page in Address Book, first time in the profile

1. Create a new profile.
2. Start Outlook.
3. Create a new message.
4. Click **To.**
5. Scroll down one page in the address list.

ABVL2: Scroll down one page in Address Book, first session

1. Start Outlook.
2. Create a new message.
3. Click **To.**
4. Scroll down one page in the address list.

AD1: Check Properties on a name in the Address Book, first Outlook session

1. Start Outlook.
2. Create a new message.
3. Click **To.**
4. Select an address.
5. Click **Properties.**

AD2: Double-click name in To line, first Outlook session

1. Start Outlook.
2. Create a new message.
3. Enter the alias selected in test AD1.
4. Click **Check Names.**
5. Double-click name.

AD3: Check Properties on a name in the Address Book, second Outlook session

This test preparation was performed after AD1

1. Create a new message.
2. Click **To.**
3. Select an address.
4. Click **Properties.**

AD4: Double-click name in To line, second Outlook session

This test preparation was performed after AD2.

1. Create a new message.
2. Enter the alias selected in test AD1
3. Click **Check Names.**
4. Double-click name.

Outlook 2000 and Outlook 97 Results

For the tabulated results of the directory access tests for each mail client, see DirectoryAccess.xls on the *Exchange* 2000 *Server Resource Kit* companion *CD*.

Outlook 2000 and Outlook 97 Measurements Analysis

A first analysis shows that Outlook 2000 queries the global catalog server only to resolve names. Only upon the second launch of Outlook 2000 (using the profile) is the global catalog server queried. When a new profile is used, the GAL name is not known, so Outlook cannot resolve it. It is during profile initialization (the first log on) that Exchange 2000 gives the client the GAL name. Outlook 2000 stores this name in its profile in a registry entry.

Upon the second launch, every initial name resolution (manual or automatic) in the Outlook 2000 session generates traffic with the Exchange server and the global catalog server. Subsequent queries are sent directly to the global catalog server. The Exchange 2000 server does not process any more name queries.

On the other hand, Outlook 97 communicates only with the Exchange 2000 server, which functions like a proxy server between the global catalog server and Outlook 97.

Outlook 2000 caches all resolved names during the session. This cache is only written or read when users click **Check Names** (or CTRL+K). The automatic resolution does not use cache; it uses the global catalog. The automatic resolution always contacts the Exchange 2000 server and the global catalog server during the first resolution in the session.

The ambiguous-name resolution generates the same amount of traffic on Outlook 97 and Outlook 2000. ANR always generates the same levels of traffic with **Check Names,** but it generates less traffic when using automatic resolution on a resolved name.

The traffic generated when the global address list is displayed in the address book does not depend on the number of entries. The initial traffic only depends on the visible entries on the first page.

In Outlook 2000, viewing address details from the **To** box generates 70 percent more bytes than viewing details from the address book. In Outlook 97, both methods generate the same amount of traffic. An initialization phase occurs when viewing details during the first Outlook session, because the details form design downloads. Showing other tabs during the first Outlook session also generates extra traffic.

Generally, both versions of Outlook generate about the same amount of traffic. The difference is in how the traffic is spread out. Outlook 2000 directly queries the global catalog server for all address book-related actions; this consumes fewer CPU cycles on the Exchange 2000 server.

Outlook Express: LDAP Mode

Outlook Express can resolve names with LDAP. Because LDAP is a protocol used for accessing directories, it is independent of the mailbox access protocols. You will need to create an LDAP account. It is not necessary to perform tests with POP3 and IMAP4. There is no difference between **Check Names** and **Find Address;** the same LDAP query occurs.

Test Details

The following tests were performed.

AR1: Find people

1. Open the **Find People** form.
2. Enter a non-ambiguous alias.
3. Click **Find Now.**

ANR1: Find people, two names returned

1. Show the **Find People** form.
2. Enter an ambiguous alias that generates two results.
3. Click **Find Now.**

ANR2: Find people, five names returned

1. Open the **Find People** form.
2. Enter an ambiguous alias that generates five results.
3. Click **Find Now.**

Outlook Express LDAP Results

For the tabulated results of the directory access tests for each mail client, see DirectoryAccess.xls on the *Exchange* 2000 *Sewer Resource Kit* companion *CD*.

Netscape Messenger: LDAP Mode

An unexpected authentication issue prevented completion of LDAP tests with Netscape Messenger.

Outlook Web Access

The only way to resolve names in Outlook Web Access is to click **Check Names** or press **CTRL+K**. The automatic resolution functionality does not exist in Outlook Web Access. Outlook Web Access does not store resolved names in a cache, so many of the tests for other clients are not useful because for Outlook Web Access the tests always produce the same result. Outlook Web Access does not provide address book searches.

Test Details

The following tests were performed.

AR1: Check Names

1. Create a new message.
2. In the **To** box, enter a non-ambiguous alias.
3. Click **Check Names.**

ANR1: Check Names, two names returned

1. Create a new message.
2. In the **To** box, enter an ambiguous alias that will generate two results.
3. Click **Check Names,** and then choose a name.

ANR2: Check Names, five names returned

1. Create a new message.
2. In the **To** box, enter an ambiguous alias that will generate five results.
3. Click **Check Names,** and then choose a name.

AL1: To button, second launch

1. Create a new message.
2. Click **To.**

AD1: To button, double-click name for details, then cancel

1. Create a new message.
2. In the **To** box, enter a non-ambiguous alias.
3. Click **Check Names.**
4. Double-click the resolved name, and then click **Cancel.**

AD2: To button, double-click name for details then OK

1. Create a new message.
2. In the **To** box, enter a non-ambiguous alias.
3. Click **Check Names.**
4. Double-click the resolved name, and then click OK.

Outlook Web Access Results

For the tabulated results of the directory access tests for each mail client, see DirectoryAccess.xls on the *Exchange* 2000 *Sewer Resource Kit* companion CD.

Outlook Web Access Measurements Analysis

Most of the Outlook Web Access traffic comes from the HTML page description and the XML content. The address-resolution process downloads only the XML content; this dramatically reduces the traffic, because the entire page is not downloaded. The traffic for checking names is related to the amount of text in the message body: when you check names the names are underlined and the Web page is refreshed. Therefore, to reduce network traffic, you can click **Check Names** before typing text in the message body or not click **Check Names** if you know that the name is unique. You can also click **Cancel** on any window with settings that you do not change.

Mail Items

These generic mail-item tests include the more common tasks performed by a messaging client: sending. reading, modifying, and deleting messages, and opening attachments.

Tests Performed

The following list describes the general tasks that each e-mail client performed for the tests.

- **Send a 1 KB, 2 KB, and 4 KB item** Send a mail message containing either 1 KB, 2 KB, or 4 KB of text. Include a simple subject indicating the contents of the message. The recipient for the message is the current logged-on user. Send messages as:
  - RTF
  - HTML
  - Plain Text
- **Send an item with a 10 KB, 50 KB, 100 KB, 500 KB, or 1,000 KB attachment** Send a mail message containing a 10 KB, 50 KB, 100 KB, 500 KB, or 1,000 KB attachment, with no other text. Include a simple subject indicating the contents of the message, such as "1 KB Attachment." Send messages as:
  - RTF
  - HTML
  - Plain Text
- **Read an item: x KB (with or without attachment)** Open the items sent in previous test.
- **Open the attachment within items** Open the attachments in the items sent in previous test.
- **Delete an item x KB (with or without attachment)** Delete the items sent in previous test.
- **Delete y items of x KB item** Empty the deleted items folder, which contains y number of x–KB sized items.
- **Process read receipt** Process a 1–KB mail message sent with a read receipt. The user double-clicks on the read receipt.

  Process delivery receipt Processed a 1–KB mail message sent with a read receipt. The user double-clicks on the read receipt.
- **Create subfolder** Create a single subfolder in the Inbox folder.

Test Details

The following tests were performed.

MSG1: New message, first Outlook session

1. Start Outlook.
2. Create a new mail message.

MSG2: New message, second Outlook session

- Create a new mail message.

SR1, SR2, SR3: Send an RTF message with 1 KB, 2 KB, or 4 KB of text

1. Open the x KB .msg file where x is 1 KB, 2 KB, or 4 KB.
2. Copy the text.
3. Create a new RTF mail message.
4. In the **To** box, enter a recipient other than yourself.
5. In the **Subject** box, enter a brief description of the test.
6. Paste the text in the message body.
7. Click **Send.**

SR4, SR5, SR6, SR7, **SR8:** Send an RTF message with 1 KB of text and an attachment of 10 KB, 50 KB, 100 KB, 500 KB, or 1,000 KB

1. Open the 1–KB .msg file.
2. Copy the text.
3. Create an RTF mail message.
4. In the **To** box, enter a recipient other than yourself.
5. In the **Subject** box, type **R1–KB***x* where x is 0050, 0100,0500, or 1000 (corresponding to the size of the attachment).
6. Paste the text in the message body.
7. Insert the corresponding attachment.
8. Click **Send.**

SH1, SH2, SH3: Send an HTML message with **1** KB, 2 KB, **or** 4 KB **of** text

1. Open the *x* KB .msg file where *x* is *1* KB, 2 KB, or 4 KB.
2. Copy the text.
3. Create a new HTML mail message.
4. In the **To** box, enter a recipient other than yourself.
5. In the **Subject** box, enter **H***x* KBOOOO where x is 1 KB, 2 KB, or 4 KB.
6. Paste the text in the message body.
7. Click **Send.**

SH4, SH5, SH6, SH7, SH8: Send an HTML message with **1** KB **of** text and an attachment **of** 10 KB, 50 KB, 100 KB, 500 KB, **or** 1,000 KB

1. Open the *1–KB* .msg file.
2. Copy the text.
3. Create a new HTML mail message.
4. In the **To** box, enter a recipient other than yourself.
5. In the **Subject** box, type **H1 KB***x* where x is 0010, 0050, 0100, 0500, or 1000 (corresponding to the size of the attachment).
6. Paste the text in the message body.
7. Insert the corresponding attachment.
8. Click **Send.**

SP1, SP2, SP3: Send a Plain Text message with **1** KB, 2 KB, **or** 4 KB **of** text

1. Open the *x–KB* .msg file where x is 1 KB, 2 KB, or 4 KB.
2. Copy the text.
3. Create a new Plain Text mail message.
4. In the **To** box, enter a recipient other than yourself.
5. In the **Subject** box, type **P***x* KBOOOO where x is 1 KB, 2 KB, or 4 KB.
6. Paste the text in the message body.
7. Click **Send.**

SP4, **SP5,** SP6, SP7, **SP8:** Send a Plain Text message with 1 KB **of** text and an attachment **of** 10 KB, 50 KB, 100 KB, 500 KB, **or** 1,000 KB

1. Open the 1KB .msg file.
2. Copy the text.
3. Create a new Plain Text mail message.
4. In the **To** box, enter a recipient other than yourself.
5. In the **Subject** box, type P1 KB*x* where x is 0050,0100,0500, or 1000 (corresponding to the size of the attachment).
6. Paste the text in the message body.
7. Insert the corresponding attachment.
8. Click **Send.**

RR1, RR2, RR3: Open an RTF message with 1 KB, 2 KB, **or** 4 KB **of** text

- Open the RTF message R*x* KBOOOO where x is 1 KB, 2 KB, or 4 KB.

RR4: Open an RTF message with 1 KB **of** text and an attachment **of** 10 KB

- Open the RTF message R1 KBOOlO and open the attachment.

RH1, RH2, RH3: Open an HTML message with 1 KB, 2 KB, or 4 KB **of** text

- Open the HTML message H*x* KBOOOO where x is 1 KB, *2* KB, or 4 KB.

RH4: Open an HTML message with 1 KB **of** text and an attachment **of** 10 KB

- Open the HTML message H1 KBOOlO and open the attachment.

RP1, RP2, RP3: Open a Plain Text message with 1 KB, 2 KB, **or** 4 KB **of** text

- Open the Plain Text message Px KBOOOO where x is 1 KB, *2* KB, or 4 KB.

RP4: Open a Plain Text message with 1 KB **of** text and an attachment **of** 10 KB

- Open the Plain Text message P1 KBOOlO and open the attachment.

OR4, OR5, OR6, OR7, OR8: Open the 10-KB, 50-KB, 100-KB, 500-KB, and 1,000-KB attachments in the 1-KB RTF messages

1. Open the RTF message R1 KBOOlO without opening the attachment.
2. Open the attachment.

OH4, OH5, OH6, OH7, OH8: Open the 10-KB, 50-KB, 100-KB, 500-KB, and 1,000-KB attachment in the 1-KB HTML message

1. Open the HTML message H1 KB0010 without opening the attachment.
2. Open the attachment.

OP4, OP5, OP6, OP7, OP8: Open the 10-KB, 50-KB, 100-KB, 500-KB, and 1,000-KB attachments in the **1-KB Plain** Text messages

1. Open the Plain Text message, P1 KB0010, without opening the attachment.
2. Open the attachment

D1: Delete a message with no attachments

1. Select a message with no attachments.
2. Press **Delete.**

D2: **Delete** a message with one attachment

1. Select a message with one attachment.
2. Press **Delete.**

PRR1: Open a read receipt

- Open a read receipt.

PDR1: Open a **delivery** receipt

- Open a delivery receipt.

DI1: Empty **Deleted Items** containing a 1-KB message

1. Put one 1–KB message in the **Deleted Items** folder.
2. Empty the **Deleted Items** folder.

DI2: Empty **Deleted Items** containing one 1-KB message with one 1,000-KB attachment

1. Put one 1–KB message containing one 1,000–KB attachment in the **Deleted Items** folder.
2. Empty the **Deleted Items** folder.

DI3: Empty **Deleted Items** containing ten **1KB** messages

1. Put ten 1–KB messages in the **Deleted Items** folder.
2. Empty the **Deleted Items** folder.

DI4: Empty Deleted Items containing ten 1-KB messages with one 1,000-KB attachment each

1. Put ten 1–KB messages each containing a 1,000–KB attachment in the **Deleted Items** folder.
2. Empty the **Deleted Items** folder.

FOL1: Create one folder

1. Open Inbox.
2. Create a new folder.

Outlook 2000 Results

For the tabulated results of the mail item tests for each mail client, see Mail.xls on the *Exchange 2000 Sewer Resource Kit* companion *CD*.

With Outlook 2000, an attachment is loaded onto the server as soon as it is inserted in the message. When an attachment is inserted in a message before the user composes the message, the attachment has more time to load on the server and, when the user clicks **Send,** the time before the message is sent is reduced.

This functionality is dependant attachment size. A 50–KB attachment immediately uploads during composition, while 10–KB attachments upload only when the user clicks **Send.**

The amount of traffic generated when a new folder is created depends on the number of folders in the mailbox. In these tests, the capture occurred with standard Outlook folders. Every new folder adds about 100 bytes in generated traffic.

The deletion (tests D1, D2) of items moves the items to the **Deleted Items** folder

HTML messages consume more bandwidth than standard Outlook messages. However, the amount of bandwidth consumed is dependent on the content of the message. The tests on public folders later in this chapter show that message form can dramatically reduce or increase the size difference between HTML format, RTF, or Plain Text. HTML is the most voluminous message format, and RTF is the lightest format in Outlook 2000. Outlook 2000 does not download the attachment, if there is one, when opening a message.

The size difference between message types is mainly due to attachment size. There is no preferred format in which to send an attachment, because attachment conversion does not affect the message size: 1 MB stays at about 1 MB after conversion. The size difference between formats is due to the 1 KB of converted text.

Outlook 97 Results

For the tabulated results of the mail item tests for each mail client, see Mail.xls on the *Exchange 2000 Server Resource Kit* companion *CD*.

Outlook 97 can only send RTF messages. The **Deleting an item** option means that messages are moved to the **Deleted Items** folder. The act of emptying the **Deleted Items** folder makes no impact on network traffic, whether or not messages have attachments.

Outlook Express IMAP Results

For the tabulated results of the mail item tests for each mail client, see Mail.xls on the *Exchange 2000 Server Resource Kit* companion *CD*.

By default, the **Save copy of sent message in the Sent Items Folder** option is selected. This option causes the message to be sent twice. When you clear it, you can divide all sending operations by two.

When opening a message with attachments, the message and the attachments are stored locally. If they are both read later, there is no more traffic. Netscape Messenger generates traffic each time you want to read the attachment.

As soon as a message is received on the server, its header is sent to the client. There is no need to generate a refresh.

The conclusions are the same as those with Outlook 2000. There is a big difference in message size between a message sent in HTLM format and a message sent in Plain Text.

The small difference in message size between formats is about the text format itself. There is no correlation between the attachment size and the difference in message size.

Outlook Express POP Results

For the tabulated results of the mail item tests for each mail client, see Mail.xls on the *Exchange 2000 Sewer Resource Kit* companion *CD*.

There is no refresh test in POP3 mode; the amount of network traffic is the same as that for reading messages.

HTML format consumes more bandwidth than Plain Text, even if the text has no rich-text formatting.

The difference between POP and HTML formats is mainly caused by the size of attachments. There is no preferred format in which to send an attachment. Attachment conversion does not change the amount of network traffic: 1 MB stays at about I MB after conversion. The size difference between formats is due to the 1 KB of converted text.

Netscape Messenger IMAP Results

For the tabulated results of the mail item tests for each mail client, see Mail.xls on the *Exchange 2000 Server Resource Kit* companion *CD*.

By default, Netscape Messenger stores sent items locally in the **Sent** folder. Outlook Express stores sent items on a server folder. To ensure equal testing conditions, the Netscape **Sent** folder was moved to the server.

Netscape Messenger does not download attachments with messages. Netscape Messenger downloads only message bodies and then downloads the attachments when users want access to them. The network traffic test captures at this point included attachment download.

If a user closes the attachment and at a later point wants to access the attachment again, Netscape Messenger downloads it again; it does not mark the attachment as already available locally.

HTML messages consume more bandwidth than Plain Text messages. This is true for both Netscape Messenger and Outlook Express.

The difference in bandwidth consumption is mainly due to attachment size. There is no preferred format in which to send an attachment. Attachment conversion does not increase the difference.

Netscape Messenger POP Results

For the tabulated results of the mail item tests for each mail client, see Mail.xls on the *Exchangr 2000 Server Resource Kit* companion *CD*.

For Netscape in IMAP mode, there is no difference in sending in HTML or Plain Text.

There were similar results with or without attachments. It made no difference in traffic consumption for different formats when no rich characters were detected.

HTML messages consume more bandwidth than Plain Text messages, for both Netscape Messenger and Outlook Express.

The difference in bandwidth used to send these message types is mainly due to attachment size. There is no preferred format in which to send attachments. Attachment conversion does not significantly increase the difference.

Outlook Web Access

For all the below tests, the default Inbox view was the **Messages** view.

Test Details

The following tests were performed:

MSG1: New Message, first time in the Windows profile

1. Log on to Windows.
2. Run Outlook Web Access (http://LONDON-04/exchange).
3. In Inbox, click **New** on the Outlook toolbar.

MSG2: Attachments

- In a new message, click **Attachments.**

REFO, 1, 2, 3: Check for new messages when 0, 1, 5 or 10 new messages

1. Send x new message(s) to the mailbox where x is 0 (1, 5, 10).
2. Click Check for new messages.

SH1, 2, 3: Send a message with 1 KB, 2 KB, and 4 KB of text

1. Open the *x* KB reference .txt file where x is 1 (2, 4).
2. Copy the text.
3. Create a new message.
4. In the **To** field, enter a recipient other than yourself.
5. In the **Subject** field, enter **Hx** KBOOOO where x is 1 (2, 4).
6. Paste the text in the main message body.
7. Click **Send.**

SH4, 5, 6, 7, 8: Send a message with 1 KB of text and an attachment of **10, 50, 100, 500,** and 1000 KB

1. Open the x KB reference .txt file where x is 1 (2, 4).
2. Copy the text.
3. Create a new HTML mail message.
4. In the To field, enter a recipient other than yourself.
5. In the **Subject** field, enter **H1 KB*x*** where x is 0010 (0050,0100,0500, 1000).
6. Paste the text in the main message body.
7. Insert the corresponding attachment.
8. Click **Send.**

RH1, 2, 3: Open a message with 1 KB, 2 KB, and 4 KB of text

- Open the message, Hx KB0000, where x is 1 (2, 4).

RH4: Open a message with 1 KB of text and an attachment of any size

- Open the message, H1 KB0010 without opening the attachment.

OH4, 5, 6, 7, 8: Open the 10 KB, 50 KB, 100 KB, 500 KB, and **1000** KB attachment in the 1 KB message

1. Open the message, H1 KB0010, without opening the attachment.
2. Open the attachment.

D1: Delete a message (with no attachment)

1. Select a message among many messages in the current folder with no attachment.
2. Press **Delete.**

D2: **Delete** a message (with one attachment)

1. Select a message among many messages in the current folder with one attachment.
2. Press **Delete.**

D3: **Delete** the **last** message (with no attachment)

1. Select the last message in the current folder with no attachment.
2. Press **Delete.**

D4: Delete the last message (with one attachment)

1. Select the last message in the current folder with one attachment.
2. Press **Delete.**

PRR1: Open a read receipt

- Open a read receipt.

PDR1: Open a delivery receipt

- Open a delivery receipt.

DI1: Empty Deleted Items containing one **1KB** message

1. Put one 1 KB message in **Deleted Items.**
2. Empty Deleted Items.

DI2: Empty Deleted Items containing one **1KB** message with one 1000 **KB** attachment

1. Put one 1 KB message containing one 1000 KB attachment in **Deleted Items.**
2. Empty Deleted Items.

DI3: Empty Deleted Items containing ten **1KB** messages

1. Put ten 1 KB messages in **Deleted Items.**
2. Empty Deleted Items.

D14: Empty Deleted Items containing ten **1KB** messages with one 1000 **KB** attachment each

1. Put ten 1 KB messages each containing a 1000 KB attachment in **Deleted Items.**
2. Empty Deleted Items.

FOL1: Create Folder (only show dialog)

- Click **Folder** in the **New** menu.

FOL2: Create Folder (full)

- Click **Folder** in the **New** menu, type **Folderl,** and then click OK.

Outlook Web Access Results

For the tabulated results of the mail item tests for each mail client, see Mail.xls on the Exchange 2000 *Sewer Resource Kit* companion *CD*.

Choosing the New Message form generates traffic only the first time you open the form. The network traffic that occurs when you delete items in a folder depends on the number of items that remain in the folder. When you delete items, the folder content is refreshed. The fewer messages that remain, the less traffic is generated. However, more traffic occurs when you delete the last item in the folder. Network traffic for deleted items does not depend on the item's size.

When you click **Attachments,** the check names action occurs, generating additional traffic for a large message body. Message attachments transfer to the Outlook Web Access server as soon as you click **Attach** in the **Attachments** dialog box. Message attachments download only when you request them.

You generate less network traffic when you click **Check for New Messages** than when you click **Refresh** on the browser toolbar.

Mail Item Analysis

HTML is the more–consuming bandwidth format, regardless of the client.

Forwarding messages does not always download attachments from the original message. It depends on the format of the original message. If the message format is Rich Text, the attachment is downloaded when the message forwards. If the format is HTML or Plain Text, the attachment is never downloaded; it stays on the server.

For instance, if you receive a plain text message with a 1–MB document, and you forward this mail, you will not download the 1–MB document from the server. This functionality is specific to Outlook 2000 because it handles messages as RTF format, and downloads attachments when you forward a message with attachments.

In POP3 mode, attachments fully download as soon as you open or read the message. This characteristic is essential for POP3 protocol. So, forwarding is similar to sending a new message.

In IMAP mode, attachment management depends on the client. Outlook Express **5.x** always downloads the whole message, including attachments. Netscape Messenger does not. In tests, Netscape Messenger fully downloaded messages with 10–KB attachments when the message was forwarded, whereas it only downloads 50–KB (and greater) attachments when you want to open the attachment. Large attachments download immediately when you click **Forward.**

There is no global difference in the way Outlook 2000 or Outlook 97 handles RTF messages. Nor is there any essential difference between the two MAPI clients.

Outlook Express is the most bandwidth consuming among all the tested clients when messages are sent in HTML format.

IMAP clients consume more bandwidth than any other client—about twice the traffic consumed by POP3 clients. Outlook Web Access generates the least network traffic for HTML-formatted messages. Therefore, Outlook Web Access is the best client for reading large text message.

The difference between Outlook Express and the other clients is not so visible because attachments create most of the traffic. The difference between IMAP and POP clients is still there. MAPI clients and Outlook Web Access are very close; they are the lowest consuming clients for HTML messages.

The difference in network traffic between IMAP and POP clients is clearly visible. The MAPI client is the lowest consumer of bandwidth.

Calendaring, Contacts and Tasks

These tests provide traffic measurements generated by Calendar. Contacts, and Tasks items. These tests were performed on the following clients:

- Microsoft Outlook 2000
- Microsoft Outlook 97
- Microsoft Outlook Web Access (no Tasks items)

Tests Performed

The following list describes the general tasks that each mail client performed for the tests

- **Open calendar** Initial view of the calendar, initiated by clicking the Calendar icon.
- **Open calendar appointment** Initial view of a new appointment, initiated by clicking **New** in the Calendar view.
- **Add a calendar item** Create a four-hour meeting in the Calendar (from 8 a.m. to 12 p.m. any day) and type **Meeting** as the description.
- **Modify a calendar item** Move the above meeting to a new time slot (12 p.m. to 4 p.m.) on the same day; rename the meeting **Another Meeting.**
- **Delete a calendar item** Delete the above appointment.
- **Open contacts** Initial view of contacts, initiated by clicking the **Contacts** item.

Open contact form Initial view of the Contact form, initiated by clicking **New** in the Contact view.

Add a contact item Create a contact with full name, company name, two telephone numbers, business address, and an e-mail address. Click **Save and Close.**

- Modify a contact item Rename the contact.
- Delete a contact item Delete the contact created above.
- Open tasks Initial view of tasks, initiated by clicking the **Tasks** item.
- Open task form Initial view of the task form, initiated by clicking **New** in the Task view.
- Add a task item Create a task called "Task".
- Modify a task item Rename it **Another task.** Click **Save and Close.**
- Delete a task item Delete the above task.

Test Details

The following tests were performed.

CAL1: Open Calendar folder

1. Start Outlook.
2. Click **Calendar** in Outlook bar.

CAL2: New appointment: first Outlook session

1. Start Outlook.
2. Click **Calendar** in Outlook bar.
3. Create a new appointment.

CAL3: Add an appointment

1. Click **Calendar** in Outlook bar.
2. Create a new appointment.
3. Fill required boxes, and then click **Save and Close.**

CAL4: Modify an appointment

1. Click **Calendar** in Outlook bar.
2. Open an appointment, modify the subject, and then click **Save and Close.**

CAL5: Delete an appointment

1. Click **Calendar** in Outlook bar.
2. Select an appointment.
3. Delete the appointment.

CTC1: Open Contacts folder

1. Start Outlook.
2. Click **Contacts** in Outlook bar.

CTC2: New contact: first Outlook session

1. Start Outlook.
2. Click **Contacts** in Outlook bar.
3. Create a new contact.

CTC3: Add a contact

1. Click **Contacts** in Outlook bar.
2. Create a new contact.
3. Fill required boxes, and then click **Save and Close.**

CTC4: Modify a contact

1. Click **Contacts** in Outlook bar.
2. Open a contact, and then modify the company.
3. Click **Save and Close.**

CTC5: Delete a contact

1. Click **Contacts** in Outlook bar.
2. Select a contact.
3. Delete the contact.

TSK1: Open Tasks folder

1. Start Outlook.
2. Click **Tasks** in Outlook bar.

TSK2: New Task: first Outlook session

1. Start Outlook.
2. Click **Tasks** in Outlook bar.
3. Create a new Task.

TSK3: Add a Task

1. Click **Tasks** in Outlook bar.
2. Create a new Task.
3. Fill required boxes, and then click **Save and Close.**

TSK4: Modify a Task

1. Click **Tasks** in Outlook bar.
2. Open a Task, modify the subject, and then click **Save and Close.**

TSK5: Delete a Task

1. Click **Tasks** in Outlook bar.
2. Select a task.
3. Delete the task

Outlook 2000, Outlook 97, and Outlook Web Access Results

For the tabulated results of the calendaring, contacts, and tasks tests for each mail client, see CalendaringContactsTasks.xls on the *Exchange* 2000 *Server Resource Kit* companion CD.

Calendaring, Contacts, and Tasks Analysis

Calendar actions are very similar with both MAPI clients, except when opening the calendar folder where Outlook 2000 has more features than Outlook 97. Outlook Web Access is more bandwidth–consuming. This difference is apparent when you click a folder or open an item. Outlook Web Access downloads a new form only once during a session and stores the form in the cache. This is similar to MAPI clients, which retrieve the form from locally-stored templates.

Similar conclusions can be drawn for contact functionality. Outlook Web Access downloads more data than MAPI clients to display the contacts forms.

Outlook 2000 and Outlook 97 are equivalent in tasks operations and related network consumption.

Public Folders

The public folder tests provide traffic measurements generated by public folders. They include MAPI and NNTP access. The tests were perfonned on the following clients.

- Microsoft Outlook 2000
- Microsoft Outlook 97
- Microsoft Outlook Express 5.0 (with an NNTP account)
- Netscape Messenger 4.7 (with an NNTP account)
- Microsoft Outlook Web Access

For MAPI-client testing, mail was not sent to a public folder' instead, the New Post form was used

Tests Performed

The following list describes the general tasks that each mail client performed for the tests.

- **Connection** When the user clicks **Public Folders,** an initial connection occurs with the associated public folders server.
- **Hierarchy listing** When the user double-clicks **All Public Folders,** the hierarchy of public folders, within the given organization appears.
- **Post an item x KB (1 KB, 2 KB, 4 KB)** Post a message containing x KB of text. Include a simple subject indicating the contents of the message, such as "*x* KB of text." Create posts using the following formats.
  - RTF
  - HTML
  - Plain Text
- **Post an item with attachment of x KB (10 KB, 50 KB, 100 KB, 500 KB, 1,000 KB)** Post a message containing an *x* KB attachment with no other text. Include a simple subject indicating the contents of the message, such as "x KB Attachment". Create posts using the following formats.
  - RTF
  - HTML
  - Plain Text
- **Read an item: x KB (with or without attachment)** Open the items posted in the previous two tests.

- Open the attachment within items Open the attachments posted in the earlier test.
- Delete an item with or without attachment Delete the items posted in the earlier test.

Test Details

The test steps are identical to the tests for generic mail items.

In addition, the following tests were also run.

CNX1: Connection, first Outlook session

1. Start Outlook.
2. Show **Folder** view.
3. Double-click Public Folders.

CNX2: Connection, second Outlook session

- Double-click Public Folders.

PF1, PF5, PF10: Open All Public Folders, with **1, 5,** or 10 public **folder(s),** first Outlook session

1. Start Outlook.
2. Show Folder view.
3. Double-click Public Folders.
4. Double-click All Public Folders containing 1, 5, or 10 public folder(s).

PF2, PF6, **PF11:** Open All Public Folders, with **1, 5,** or 10 public folders, second Outlook session

1. Start Outlook.
2. Show Folder view.
3. Double-click Public Folders.
4. Double-click All Public Folders containing 1, 5, or 10 public folder(s).

The following tests were only run on NNTP clients.

PF1, PF5, PF10: Refresh Newsgroups List, with **1, 5,** or 10 public folders, first client session

1. Start the NNTP client.
2. Right-click the newsgroups server, and then choose Newsgroups.
3. Refresh the newsgroups list (In Outlook Express, click Reset List).

PF2, PF6, PF11: Open All Public Folders, with 1, 5, or 10 public folders, second client session

This test was run after the PF1, PF5, and PF10 test.

1. Right-click the newsgroups server, and then choose **Newsgroups.**
2. Refresh the newsgroups list (In Outlook Express, click **Reset List).**

Outlook 2000 Results

For the tabulated results of the public folder tests for each mail client, see PublicFolders.xls on the *Exchange* 2000 *Sewer Resource Kit* companion *CD*.

After the initial connection from the Outlook client to the Public Folder server, each subsequent connection (in the same MAPI session) generated traffic. Clicking **Public Folders** does not generate traffic based on the total number of public folders. Instead, the traffic generated is dependent on the number of folders within the **Favorites** folder. The status of the folders in the **Favorites** folder updates to reflect the number of unread messages.

Clicking **All Public Folders** generates traffic based on the number of first-level folders; sub folders do not generate traffic at this point. The first connection to **All Public Folders** generates more traffic than subsequent connections, and these only check for changes within the hierarchy.

For generic mail items, deleting a publication item is not effected by the message size; rather, traffic levels rely on the existence of attachments. The traffic is equal to one or many attachments.

HTML posts are larger than any other formats. In HTML tests, the text used was not the same as in the generic mail item tests; instead, text from Readme files were used. HTML text conversion depends on the form of the text format: the more line breaks, the bigger the converted text.

Message text formats do not impact network traffic when attachments are inserted. The small difference with 100 KB in HTML can be due to an upload, which begins when a message posts.

Outlook 97 Results

For the tabulated results of the public folder tests for each mail client, see PublicFolders.xls on the *Exchange* 2000 *Server Resource Kit* companion CD.

Outlook Express NNTP Results

For the tabulated results of the public folder tests for each mail client, see PublicFolders.xls on the *Exchange* 2000 *Server Resource Kit* companion CD.

In general, less traffic is generated after the initial connection. During the initial connection, the client requests and receives appropriate permissions, connects to the NNTP server, and downloads the newsgroups list. The client keeps the list in a local cache. In Outlook Express, the newsgroup list only shows the newsgroup name and its description.

HTML publications are still larger than Plain Text publications; however, sending attachments is not format–sensitive.

Netscape Messenger NNTP Results

For the tabulated results of the public folder tests for each mail client, see PublicFolders.xls on the *Exchange* 2000 *Server Resource Kit* companion CD.

Usually, less traffic is generated after the initial connection. During the initial connection, the client requests and receives appropriate permissions, connects to the NNTP server, and downloads the newsgroups list. The client keeps the list in a local cache. With Netscape Messenger, the list shows the newsgroup name, its description, and the number of unread posts. This explains the difference in network traffic consumed by Outlook Express and Netscape.

Outlook Web Access Results

Public folders are available only on the **Folders** tab in the Outlook bar. Public folder favorites do not exist; thus, there is no All Public Folders root folder: **Public folders** contains the entire public folder hierarchy.

For the tabulated results of the public folder tests for each mail client, see PublicFolders.xls on the *Exchange* 2000 *Server Resource Kit* companion CD.

The whole hierarchy (first level of Public Folders) is downloaded when the user clicks the **Folders** on the Outlook bar. Thus, no further traffic is generated when the user opens the public folders tree or clicks **Folders** again.

Deleting a public folder item produces the same network traffic as deleting a message item

Public Folder Analysis

Once connection and hierarchy listings have been established, less traffic occurs during similar queries in the same session.

Outlook Express generates more traffic than Netscape Messenger in posting HTML publications, regardless of the size of the messages being handled. However, there is a small difference in network traffic generated by Outlook Express and Netscape for posting HTML publications with attachments. Outlook Express needs, on average, 2 KB more bandwidth than Netscape to get a publication with an attachment.

The difference in network traffic is proportional to attachment size. With a 50–KB attachment, there is difference of 1 KB. While with a 1–MB attachment, there is gap of 36 KB. Outlook Express generates more traffic than Netscape when downloading plain text publicaiions with attachments.

Outlook 2000 with Terminal Services

Terminal Services, running on a computer running Windows 2000, enables all client application execution, data processing, and storage to occur on the server. Applications appear on the user's device by means of terminal emulator software, which can run on a variety of client hardware devices: a personal computer, a Windows-based terminal, or even a Windows CE-based handheld device. The term *Windows-based terminal* broadly refers to a class of thin client terminal devices, such as Terminal Services, that can gain access to servers running a multi-user Windows operating system.

Users can gain access to Terminal Services over any TCP/IP connection, including Remote Access Ethernet, the Internet, wireless network, WAN, or virtual private network (VPN). The user experience is only limited by the characteristics of the weakest link in the connection (for example, hardware capabilities), and the security of the link is governed by the TCP/IP deployment in the data center.

With Terminal Services, the terminal emulation software sends keystrokes and mouse movements to the server. Terminal Services manipulates data locally and then passes back the display. This approach allows remote control of servers and centralized application management, minimizing network bandwidth requirements between the server and client.

Some companies may choose to provide access to Outlook 2000 through terminal emulation software. This section evaluates the impact on network traffic, when using Outlook 2000 through the 32-bit Terminal Services client, on a computer running Windows 2000 Server.

Test Details

The Terminal Services client sends keystrokes and mouse movements to the server. For that reason, the tests provide differing results because it is impossible to ensure exact mouse movements. These tests were performed by moving the mouse very carefully, generating as little additional traffic as possible. The traffic captured in these tests is only that between the client and the Terminal server. During a Terminal Services session, no traffic occurs between the Terminal client computer and the Exchange server.

The following tests were performed.

Connect

Captures the traffic generated by initial connection to the server. The capture was stopped when the logon dialog box appeared.

CONNECT: open a Terminal Services session

1. Start Terminal Services client.
3. Click **Connect.**

TS logon

Captures the traffic generated by opening a session on the server. The test stopped when the desktop appeared.

TSLOGON: open a session on the Terminal sewer

1. Enter user name and password.
2. Press ENTER.

TS Logoff

Captures the traffic resulting from closing a session on the server.

TSLOGOFF: close the Terminal Services session

1. Quit all running applications.
2. Click **Start,** click **Shutdown.**
3. In the **Shut down Windows** box, select **Log off *current–user.***

Outlook Logon

Captures the traffic generated by opening an Outlook session on the Terminal server. The Outlook preview pane is open.

LOGON1: open Outlook containing Outlook welcome message

1. Create a profile in an empty mailbox.
2. Double-click Outlook icon on the desktop.

LOGON2: open Outlook containing no messages

1. Create a profile in an empty mailbox.
2. Double-click Outlook icon on the desktop.

LOGON3: open Outlook containing five messages

1. Put five unread messages in your Inbox.
2. Quit Outlook.
3. Double-click Outlook icon on the desktop.

LOGON4: open Outlook containing ten messages

1. Put ten unread messages in your Inbox.
2. Quit Outlook.
3. Double-click Outlook icon on the desktop.

Outlook Logoff

Captures the traffic resulting from closing the Outlook session.

LOGOFF: logoff Outlook

1. Close all Outlook windows, except the main window.
2. Press ALT+F4.

Message Actions

Captures the traffic resulting from the last step in each of the following tests.

MSG1: Open Inbox folder containing no messages

1. Open Outlook.
2. Open **Deleted Items** folder.
3. Click **Inbox** on the Outlook bar.

MSG2: Display the new message form

- Click **File,** point to **New,** and then click **Mail Message.**

MSG3: New message form with CTRL+N

- From the Inbox view, press CTRL+N.

MSG4: New message form with CTRL+N after x attempts

- From the Inbox view, press CTRL+N.

MSG5: Send a message

1. Click **File,** point to **New,** and then click **Mail Message.**
2. Fill the **To** and **Subject** boxes.
3. Click **Send.**

MSG6: Open Outlook Today containing one appointment and one task

1. Open Outlook.
2. Click Outlook Today on the Outlook bar.

Calendar Actions

Captures the traffic resulting from the last step in each of the following tests.

CAL1: Open Calendar folder with no entry

1. Open Outlook.
2. Click Calendar on the Outlook bar.

CAL2: New appointment form

- From the Calendar view, click **File,** point to **New,** and then click **Appointment.**

CAL3: New appointment form with CTRL+N

- From the Calendar view, press CTRL+N.

CAL4: New appointment form with CTRL+N afterx attempts

- From the Calendar view, press CTRL+N.

CAL5: Save and close an appointment

1. From the Calendar view, click **File,** point to **New,** and then click **Appointment.**
2. Fill the necessary boxes.
3. Click **Save and Close.**

Contact Actions

Captures the traffic resulting from the last step in each of the following tests.

CTC1: Open Contacts folder containing no entries

1. Open Outlook.
2. Click **Contacts** on the Outlook bar.

CTC2: Display the new contact form

- From the Contacts view, click **File,** point to **New,** and then click **Contact.**

CTC3: New Contact form with **CTRL+N**

- From the Contacts view, press CTRL+N

CTC4: New Contact form with **CTRL+N** after x attempts

- From the Contacts view, press CTRL+N.

CTC5: Save and close a contact

1. From the Contacts view, click **File,** point to **New,** and then click **Contact.**
2. Fill the necessary boxes.
3. Click **Save and Close.**

Task Actions

Captures the traffic resulting from the last step in each following test.

TSK1: Open Tasks folder containing no entries

1. Open Outlook.
2. Click **Tasks** on the Outlook bar.

TSK2: Display the new task form

- From the Tasks view, click **File,** point to **New,** and then click **Task.**

TSK3: Display the new Task form with **CTRL+N**

- From the Tasks view, press CTRL+N.

TSK4: New Task form with **CTRL+N** after x attempts

- From the Tasks view, press CTRL+N.

TSK5: Save and close a task

1. From the Tasks view, click **File,** point to **New,** and then click **Task.**
2. Fill the necessary boxes.
3. Click **Save and Close.**

Note Actions

Captures the traffic resulting from the last step in each following test.

NOT1: Open Notes folder containing no entries

1. Open Outlook.
2. Click **Notes** on the Outlook bar.

NOT2: Display the new Note form

- From the Notes view, click **File,** point to **New,** and then click **Note.**

NOT3: New Note form with CTRL+N

- From the Notes view, press CTRL+N.

NOT4: New Note form with CTRL+N after x attempts

- From the Notes view, press CTRL+N.

NOT5: Save and close a Note

1. From the Notes view, click **File,** point to **New,** and then click **Note.**
2. Fill the necessary boxes.
3. In the upper-right part of the screen, click **Close.**

Address Book Actions

Captures the traffic resulting from the last step in each of the following tests.

AR1: Check one name

1. Create a new mail message.
2. In the **To** box, enter an alias that will resolve to a single name.
3. Press CTRL+K (Check Names).

AR2: Full check one name

1. Create a new mail message.
2. In the **To** box, enter an alias that will resolve to a single name.
3. On the toolbar, click **Check Names.**

AL1: Display the address book with the To button

1. Create a new mail message.
2. Click **To.**

AB1: Add one address in To

1. Create a new mail message.
2. Click **To.**
3. Select one address, click **To,** and then click OK.

AD1: Show Details in address book

1. Create a new mail message.
2. Click **To.**
3. Select one address.
4. Click **Properties.**

Terminal Services Client Results

For the tabulated results of the Terminal Services tests for each mail client, see Outlook2000TerminalServices.xls on the ***Exchange*** 2000 ***Sewer Resource Kit*** companion CD.

Terminal Services Measurement Analysis

These tests were done with an 800 x 600 pixel resolution for the Terminal Services session. Other captures were done on the tests "Connect" and "Open Terminal Services session" with a lower resolution (640 x 480) and higher resolution (1024 x 768). These new captures had the same traffic volume (bytes and frames).

To reduce traffic during Terminal Services sessions, use keyboards shortcuts as much as possible. Mouse events are sent based on a regular frequency and increased network traffic. In addition, when configuring a Terminal Services connection, select the option to enable data compression; this will also decrease traffic.

Web Storage System

The Microsoft Web Storage System is a database for messaging, collaboration, rich document storage, and Web-enabled applications. The Web Storage System can be accessed by a wide range of client software, including:

- Microsoft Outlook 97, Outlook 98, and Outlook 2000 messaging and collaboration clients
- Outlook Express and any e-mail or newsgroup client that supports SMTP/POP3, IMAP4, or NNTP
- Microsoft FrontPage
- Microsoft Office 2000
- Windows Explorer
- Web Folders (included with Internet Explorer 5, Office 2000, and Windows 2000)
- Any Web browser
- The MS-DOS prompt
- Any 32-bit application for Windows

Upon installation, the Web Storage System is mapped to the M drive on the Exchange 2000 server, and accessed in the same way as the existing Windows file system. The administrator can share this virtual drive to give access to users who can then access their mailboxes and public folders as with any file server. You can access through Uniform Naming Convention (UNC) or a mapped drive letter.

Tests Performed

Seven test categories evaluated Web Storage System traffic:

- **Tools/map network drive** Connect to the network share and map a network drive with Windows 2000.
- **Net use** Map a network drive with the command line, NET USE.
- **Logoff** Disconnect from the network share.
- **Create folder** Create a new folder within your mailbox.
- **Copy an item: x-KB (10-KB, 50-KB, 100-KB, 500-KB, and 1,000-KB)** Copy a file into one of the mailbox folders.
- **Open an item: x-KB (10-KB, 50-KB, 100-KB, 500-KB, and 1,000-KB)** Open a file in one of the mailbox folders.
- **Delete an item: x-KB (10-KB, 50-KB, 100-KB, 500-KB, and 1,000-KB)** Delete a file in one of the mailbox folders.

Log On and Log Off and User Traffic Results

For the tabulated results of the Web Storage System tests for each mail client. see WebStorageSystem.xls on the *Exchange* 2000 *Server Resource Kit* companion CD.

Web Storage System Measurement Analysis

The user traffic results are similar to the traffic generated when you work with a standard shared directory on a file server. Exchange 2000 Server does not generate network traffic; all processing is performed on the server.

If you compare the user traffic results with sending or opening such attachments with a MAPI client, you will not see big difference. The Web Storage System is not a particularly more or less efficient way to work with your mailbox: it is just a convenient method for some users. It makes mailbox access easier because Collaboration Data Objects (CDO) is not required. You can now run batch files on your mailbox.

Instant Messaging

Instant Messaging consists of immediate, text-based messages that you can send to other users on a computer network. Unlike e-mail messages, these messages post to the other user's screen, providing the basis for new forms of collaboration. Instant Messaging has become a wide-scale communication phenomenon for Internet users and is poised to play a significant role as a business tool for organizations of all sizes. Exchange 2000 includes an Instant Messaging service built on a secure, standards-oriented architecture ideally suited for both enterprise deployment and deployments across the Internet for business-to-business communication. The client software for instant messaging in Exchange 2000 is the MSN Messenger client.

Presence information, closely related to Instant Messaging, enables one computer user to see whether another user is logged onto a network, corporate LAN, or the Internet. Exchange 2000 provides complete support for presence information.

Tests Performed

The following tests were performed.

- **Logon with 0, 1 and 10 contacts** Open an Instant Messaging session with no contacts, then 1 contact, then 10 configured contacts.
- **Add a contact** Add a contact in the contacts list.
- **Delete a contact** Delete a contact from the contacts list.
- **Change status** Statuses include:
  - Online

- Be right back
- Away
- On the phone
- Out to lunch
- Appear offline

The traffic was tracked on the contact side, and on the user side (where the contact is added).

- **Send and receive messages** This function is divided into three parts:
  - Entering characters with keyboard
  - Sending the message
  - Receiving the message

Log On and Log Off and User Traffic Results

For the tabulated results of the Instant Messaging tests for each mail client, see InstantMessaging.xls on the Exchange 2000 *Sewer Resource Kit* companion *CD*.

The Instant Messaging client generates network traffic when the status changes or when messages are sent.

Status Change

When users change their status, information goes to the Exchange 2000 server, which manages the Instant Messaging service. The server sends the status change to all users who have that user in their contacts list. The new status is then reflected in the contacts list.

All status changes do not generate the same network traffic volume. The status change to "Away" can happen automatically when no activity is detected. This feature is configured in the option **Show me as Away when I'm inactive for *x* minutes,** available in the MSN Messenger client. In that case, when no activity is detected for x minutes, the client switches to "Away" status, and notifies the server of the change. When an activity is detected again, the client switches back to online status.

Sending Messages

The amount of network traffic required for messages depends on touch speed. As a message is created, if the first pause occurs after the first character is typed, the sender exchanges 12 frames with the server, and the server exchanges seven frames with the recipient.

If the first pause occurs after least two characters are typed, 24 frames exchange between the server and the sender and 14 are exchanged between the server and the recipient.

Afterwards, multiples of 12 frames are exchanged between sender and server, and multiples of seven frames are exchanged between server and recipient. This traffic occurs until the user clicks **Send.** In actuality, the recipient receives the message content before sending occurs. This accounts for the rapidity of Instant Messaging.

Clicking **Send** creates 16 frames between the sender and the server, and 8 frames between the server and the recipient.

A copy or paste into the message body does not generate any traffic.

Client Traffic Measurement Conclusions

There are many things to learn from these tests.

- Outlook 2000 and Outlook 97 require approximately the same bandwidth.
- Logging on represents a large part of the traffic with MAPI clients.
- Outlook 2000 is more integrated with Windows 2000 because of the Outlook capacity to query Active Directory.
- Of the different formats tested, HTML messages generate the most traffic, and RTF messages generate the least.
- Attachments only download when users want to see them.
- Except for RTF messages, no compression occurs between clients and servers. The full name is RTF-compressed, so messages are compressed between MAPI clients and the server. Therefore, HTML (no compression) generates more network traffic when messages are downloaded. Many third-party tools fix this critical issue for remote users with low bandwidth connections.
- Among Internet protocols, POP3 generates the least network traffic. Netscape Messenger seems to send HTML messages more efficiently than Outlook Express.
- The use of XML in Outlook Web Access significantly reduces traffic because the browser stores forms in the cache.. In some situations, Outlook Web Access should be considered over Outlook 2000 with Terminal Services.
- Mouse movements generate excess network traffic when using Outlook 2000 with Terminal Services. Use keyboard commands to significantly reduce network traffic.
- Web folder offers a new and easy way to access to mailboxes or publics folders. Because Web folder access generates the same traffic as common file share access, it does not increase network traffic.

DSProxy

This section includes detailed information about DSProxy and client access to Active Directory.

The DSProxy process allows MAPI clients to access Active Directory. It performs two major functions:

1. Sending directory requests on behalf of clients to Active Directory through the Name Service Provider Interface (NSPI).
2. Referring smart MAPI clients directly to the Active Directory (RFR).

At startup, the Exchange System Attendant finds the most appropriate Active Directory server in the domain through the Domain Name System (DNS), then resolves and passes its name through to the DSProxy process (Dsproxy.dll). This is signaled by event 9010 for the MSExchangeSA process.

It is also possible to ascertain which Active Directory domain controller a particular Exchange server is using by means of Exchange computer properties in the Microsoft Exchange System Manager snap-in.

In some situations, the administrator can specifically set the server that DSProxy uses. You can accomplish this by changing the following registry entries:

To specify the global catalog server for earlier MAPI clients

1. On the **Run** line, type **regedit.exe**, and then click OK.
2. In the registry editor, navigate to HKEY_LOCAL_MACHINE\SYSTEM \CurrentControlSet\Services\MSExchangeSA\Parameters.
3. Select the **NSPITargetServer** entry.
4. Assign the name of the global catalog server to specify the global catalog server for earlier MAPI clients that use only NSPI.
5. Close the registry editor.

Caution Do not use a registry editor to edit the registry directly unless you have no alternative. The registry editors bypass the standard safeguards provided by administrative tools. These safeguards prevent you from entering conflicting settings or those likely to degrade performance or damage your system. Editing the registry directly can have serious, unexpected consequences that prevent the system from starting, thus requiring a re-install of Windows 2000 or Exchange 2000. To configure or customize Windows 2000, use the programs in Control Panel or Microsoft Management Consoie (MMC) whenever possible.

To specify the global catalog server for recent MAPI clients (Outlook 2000) which can use RFR

1. On the **Run** line, type **regedit.exe**, and then click OK.

   In the registry editor, navigate to HKEY_LOCAL_MACHINE\SYSTEM \CurrentControlSet\Services\MSExchangeSA\Parameters.

2. Select the **RFRTargetServer** entry.

3. Assign the name of the global catalog server to specify the global catalog server for recent MAPI clients (Outlook 2000) that use RFR.

4. Close the registry editor.

The DSProxy NSPI works by blindly forwarding MAPI directory system (DS) requests to a global catalog server. This means that the remote procedure call (RPC) packet is not opened or evaluated because this would incur a significant overhead on the Exchange 2000 server, and complicate the security structure. The process begins by creating a listening thread for each supported network protocol, and a single working thread for each processor. This can accommodate up to 512 client connections and dynamically adds more threads as required. A socket-mapping table keeps a reference of connections between clients and servers, ensuring that the correct responses from Active Directory pass to the associated client.

DSProxy works over TCP/IP, Internetwork Packet Exchange (IPX), and AppleTalk protocols; however, it does not work over network basic input/output system (NetBIOS).

The following sections explain how different mail clients access information stored in Active Directory.

Outlook 98 and Outlook 2000

Newer versions of the Outlook client, such as Outlook 98 and Outlook 2000 (version 8.5.6204.0 or later), use a slightly different method for address book access. Initially, Outlook expects to find the directory service on the home Exchange server. Because the version of the Exchange server running on the system can be determined only after loading Emsmdb32.dll (which is after the address book provider—Emsabp.dll), Outlook will go through the DSProxy process for the first session. Once the client has contacted the DSProxy service (it will try all available transport protocols), a special referral passes back to the client, informing that client that all future directory requests should be sent to the global catalog server. Outlook stores this referral in the MAPI profile in the registry.

The referral mechanism reduces the load on the Exchange 2000 server and address book lookup latency; however, when an explicit server name is entered into the profile, Outlook requires a restart if that Active Directory server fails. If that occurs, the Exchange 2000 server passes Outlook a new referral.

Some scenarios require Outlook clients, even the latest versions, to go through the DSProxy process without being referred. For example, when a firewall exists between client computers and Active Directory servers, the firewall can be opened up to allow the Exchange 2000 server to access Active Directory.

To prevent Exchange from returning referrals

1. On the **Run** line, type **regedit.exe**, and then click OK.
2. In the registry editor, navigate to HKEY_LOCAL_MACHINE\SYSTEM \CurrentControlSet\Services\MSExchangeSA\Parameters.
3. Click **Edit,** point to **New,** and choose **DWORD Value.**
4. Type **NoRFRService** to name the entry.
5. Right-click **NoRFRService**, click **Modify,** and then assign a value of **1** to prevent Exchange from returning directory referrals.
6. Close the registry editor.

Caution Do not use a registry editor to edit the registry directly unless you have no alternative. The registry editors bypass the standard safeguards provided by administrative tools. These safeguards prevent you from entering conflicting settings or those likely to degrade performance or damage your system. Editing the registry directly can have serious, unexpected consequences that prevent the system from starting, thus requiring a re-install of Windows 2000 or Exchange 2000. To configure or customize Windows 2000, use the programs in Control Panel or Microsoft Management Console (MMC) whenever possible.

Exchange Client and Outlook 97 and Outlook 98

Older MAPI clients, such as Exchange 97, Exchange 98, Outlook 97, and Outlook 98 (version 8.5.6204.0 or earlier), make MAP1 Directory Service (MAPI DS) requests to an Exchange server. These requests range from resolving a fragment of text typed on the **To** line of a message into a user name, to showing objects from the global address list. To make Exchange 2000 compatible with the existing MAP1 client base, an Exchange 2000 server sends any MAPI DS requests through to a local global catalog server on the network. The DSProxy process on the Exchange 2000 server accomplishes this task and forwards the packet; it does not change the request into Lightweight Directory Access Protocol (LDAP). Active Directory supports a number of protocols, including LDAP and MAPI DS, so an Outlook directory request is valid, even against an Active Directory server.

The global catalog server returns the result to the Exchange 2000 server, and then sends the result on to the MAPI client. This process is seamless to the user and takes very little time to complete.

Traffic and Load Generated Through the DSProxy Process

The following steps describe the communication process for a one-recipient name lookup:

1. The MAPI client sends one network packet to the Exchange 2000 server. The packet contains the name to look up, in plain text.
2. The Exchange 2000 server sends the request to a local global catalog server.
3. The local global catalog server returns the result to the Exchange 2000 server.
4. The Exchange 2000 server returns the result to the MAPI client.
5. The MAPI client returns an acknowledgement to the Exchange 2000 server.
6. The Exchange 2000 server sends the acknowledgement to the local global catalog server.

The directory lookup only produces six frames on the network. Even when multiple names need to be looked up in the directory, the name fragments are sent in one request packet.

If the user chooses to browse the global address list, the same process takes place. A few extra frames are seen on the network as the user scrolls through the address list, thus causing the client to retrieve more information.

Resource Guide Table of Contents

CHAPTER 26 Exchange 2000 Architecture 659

Server Architecture 659

IIS Process 661

Web Storage System 661

Exchange Interprocess Communication Layer 661

Supported Internet Protocols 664

POP3 and IMAP4 672

NNTP Support 673

Web Storage System 674

Microsoft Search Service 675

Storage Location for Documents and Applications 675

Support for Multiple Databases 675

Content Storage 676

Installable File System 677

Accessing Data using IFS 677

Enhanced Active/Active Clustering 679

Public Folders 679

Multiple Store Support 679

Using Multiple Public Folder Trees 680

Public Folders in Active Directory 681

Connecting to a Public Folder 681

CHAPTER 26 Exchange 2000 Architecture *(continued)*

Message Flow 683
Message Routing Architecture 683
Routing Group Master and Propagation 683
Link State Algorithm 684
Message Transport 684
Component Storage Location 684
Windows 2000 Message Flow 686
Inbound Message Flow from SMTP to Web Storage System 687
Outbound Message Flow From Web Storage System to SMTP 689
Inbound Message Flow Through the MTA 691
Outlook Web Access Architecture 693
Accessing a Server 693
Server Components 695
Log On and Mailbox Display 695
Logout 696
Opening an Item 696
Front-End and Back-End Servers 697
Outlook Web Access Functions 698
WebDAV 698
WebDAV Process 699
Exchange 2000 and Previous Versions of Exchange 700
Architectural Improvements in Exchange 2000 700
Web Storage System Changes 701
Clustering Improvements 701
Exchange 2000 Integration with Active Directory 702
Active Directory Connector 702
Active Directory Connector process 703
Exchange 2000 Dependencies on Windows 2000 and Active Directory 706

CHAPTER 27 Application Development 707
Overview of Web-Based Collaboration 708
Network Protocols 709
HTTP and WebDAV Extension 709
SMTP 711
NNTP 711
IMAP4 712
POP3 712
LDAP 713
API Components 714
ADO 714
ADSI 719
CDO 723
OLEDB 727
MAPI 727
Formats 727
RFC 822 727
MIME 728
HTML 728
XML 728
Microsoft Web Storage System 729
Web Storage System Schema 729
IIS Integration 729
The Exchange Server Active Directory Schema 730
Exchange 2000 Web-Based Collaboration Applications 731
Microsoft Windows 2000 Server 731
Microsoft Internet Information Services 731
Active Server Pages (ASPs) 732
Microsoft Internet Explorer 734

CHAPTER 27 Application Development *(continued)*

SMTP and NNTP Inbound and Outbound Events 734
Space Allocation 734
File Semantics 735
Microsoft Outlook 2000 735
32-bit Forms 735
Instant Switching Between Forms and Run Time 735
Advanced Fields and Views 735
Built-in Forms Modules 735
Extendable Forms 736
Outlook Forms Designer 736
Visual Basic Expression Service 736
Instant Collaboration 736
Event Support in Exchange 2000 738
Using Event Sinks in Exchange 2000 738
Web Storage System Events 738
Synchronous Events 739
Asynchronous Events 742
Registering for Web Storage System Events 743
Accessing the Item 743
SMTP Transport Events 745
SMTP Protocol Events 746
SMTP Transport Components 746
Event Bindings 746
SMTP Protocol Event Sinks 747
SMTP and NNTP Transport Event Sinks with CDO 747
Event Binding Rules 748

CHAPTER 27 Application Development *(continued)*

Workflow Components For Exchange 2000 750
Workflow 750
Creating Workflow Processes 750
Using CDO Workflow Event Sinks 751
Workflow Engine 752
Adding the Workflow System Account 754
Workflow Security 755
Application Development Tools 756
Web Storage System Viewer 756
Web Storage System Schema Designer 757
The Web Storage System Application Deployment Wizard 757
The Event Sink Template Wizard for Visual Basic 6.0 757

CHAPTER 28 Backup and Restore 759

Exchange 2000 Database Technology 760
Extensible Storage Engine Structure 760
Multiple Storage Groups and Failover 761
Databases 761
Database GUID 762
Mailbox GUID 762
Transaction Log Files 762
Log File Signature 763
Checkpoint File 763
Circular Logging 763
Checksum 764
Individual Database Backup 764
Exchange 2000 Server Clusters 764

CHAPTER 28 Backup and Restore *(continued)*

Backup Prerequisites . . . 766
Planning a Backup Strategy . . . 766
Creating and Verifying Daily Backups . . . 767
Standardizing Backup Formats . . . 767
Publishing a Maintenance Schedule . . . 767
Checking for Hot Fixes . . . 767
Deciding the Types of Data to Back Up . . . 767
Resource Requirements . . . 768
Exchange 2000 Services and Permissions . . . 768
Active Directory . . . 769
Internet Information Services . . . 769
Performance Considerations . . . 770
Capturing Configuration Data . . . 770
Performing Active Directory Backups . . . 770
Backup Categories . . . 770
Full Backup . . . 771
Copy Backup . . . 771
Incremental Backup . . . 771
Differential Backup . . . 771
Backing Up Data on a Cluster Server Node . . . 772
Backup Process . . . 772
Online Backup Process . . . 773
Offline Backup . . . 775
Backup Rotation Schedule . . . 775
Month-Week-Day . . . 775
Towers of Hanoi . . . 776
Restore Process Overview . . . 777

CHAPTER 28 Backup and Restore *(continued)*

Prerequisites for Disaster Recovery 779
Planning Restore Strategies 779
Determining Overall Performance 780
Restore Categories 781
Restoring Online Backups 781
Restoring Offline Backups 781
Restoring a Single Mailbox 782
Restoring a Single Database 783
Moving a Database to a Different Storage Group 783
Restore Constraints 783
Restoring Web Storage System to a Different Server 783
Restoring Exchange Data to a Non-Production Server 784
Database Consistency 784
Defragmenting Databases 785
Online Defragmenting 785
Offline Defragmenting 785
Database Size After Defragmentation 786
Soft Recovery 786
Log Transaction Redo 786
Using ESEUTIL 787
File Header Output 787
Transaction Log File Headers 788
Repairing Damaged Databases 788
Checking Database Integrity 788
Forcing Soft Recovery 789
Using ISINTEG 789
Reference Database and Multiple Passes 790
Running ISINTEG 790

CHAPTER 28 Backup and Restore *(continued)*

Recovering from Disasters 791
Requirements for Recovering Exchange 2000 792
Recovering an Exchange 2000 Member Server 793
Reinstalling Windows 2000 793
Restoring the System Drive 793
Restoring Windows 2000 System State 794
Run Exchange Setup in Disaster Recovery Mode 794
Recovering Databases 795
Recovering an Exchange 2000 Member Server Running Site Replication Service 795
Restoring an SRS Database 796
Recovering an Exchange 2000 Member Server Running Key Management Service 796
Restoring Key Management Service Database 797
Recovering an Exchange 2000 Cluster Server 797
Recovering a Single Server in a Cluster 798
Recovering a Lost Cluster Quorum 799
Best Practices 800
Increase Backup Tape Reliability 800
Verify and Validate Backups 801
Verifying the event 801
Verifying Data 801
Document and Archive Backups 801
Test Backups 802

CHAPTER 29 Monitoring and Maintaining 803

Exchange 2000 Monitoring Features 804

Monitoring and Status Tool 804

Notifications 804

Status 806

Diagnostic Logging 806

Protocol Logging Tool 808

Using Event Viewer to View Logs 809

Queue Viewer 810

Windows 2000 Performance Monitoring Tools 811

Performance Console 811

System Monitor 811

Performance Logs and Alerts 815

Task Manager 816

Monitoring Processes 816

Monitoring the System 816

Terminal Services Client 817

Network Monitor 817

Observing Resource Usage 817

Measuring Network Traffic 818

Network Diagnosis Tool 820

Analyzing Performance Data 820

Establishing a Baseline 821

Analyzing Results 821

Watching for Large Values 822

Including Thread Identifiers 822

Ignoring Occasional Spikes 822

Using Graphs for Reporting 822

Excluding Startup Events 822

Investigating Zero Values or Missing Data 823

Identifying and Investigating Potential Bottlenecks 823

CHAPTER 30 Security **827**
Security Risks 827
Windows 2000 Security Features 829
Active Directory 829
Security Subsystem 829
Local Security Authority 830
Groups in Active Directory 831
Access Control 836
How Access Control Works 836
Exchange 2000 Access Control Model 838
Auditing 838
How Auditing Works 838
Auditing Exchange Active Directory Objects 839
Kerberos 839
How Kerberos Works 839
Kerberos and Exchange 840
Kerberos Delegation of Authentication 840
Delegation of Authentication in Exchange 841
Certificate Services 842
Certificate Service and Active Directory 842
Encrypting File System 842
Disabling EFS for All Computers in a Windows 2000 Domain 842
EFS Limitations 843
Internet Protocol Security 844
Layer 3 Protection 844
The IPSec Model 845
TCP/IP Filtering 846
Security Configuration Tool Set 847
Security Configuration 848
Security Templates 848
Analyzing Security 850

CHAPTER 30 Security *(continued)*

Exchange 2000 Security Features 851
Key Management Service 851
Encryption 851
Hash Functions 852
Ciphers 852
Algorithms 853
Certificate Services and the Key Management Service 853
Virtual Server Security 854
Virtual Server Connection Control 854
Protocol Logging 855
Permissions 856
Predefined Permissions 856
Exchange Administration Delegation Wizard 858
Levels of Administration 858
Organization Preparation 859
Other Permissions Issues 864
Securing Client and Server Communication 865
Inbound Encryption 865
Encrypted RPC 866
Securing Your Internet Connection 867
Virus Protection 867
Physical Security 867
Firewalls 868
Dual-Homed System 868
Proxy Servers 868
Domain Name System 868
Security Updates 869

CHAPTER 31 Optimizing Exchange 2000 871

Optimizing Active Directory 871

Forests 872

Single-Forest Environment 872

Multiple-Forest Environment 872

Limitations 873

Trees 873

Domains 873

Determining the Number of Domains 874

Organizational Units 875

Trust Relationships 876

One-Way Non-Transitive Trusts 876

Two-Way Transitive Trusts 877

Limiting Trust Relationships 877

Domain Controller Roles 877

Global Catalog 878

Operations Master 879

Schema Master 879

Domain Naming Master 879

Relative Identifier Master 879

Primary Domain Controller Emulator 880

Infrastructure Master 880

Windows 2000 Sites 880

Active Directory Schema 881

Schema Objects 881

Classes 881

Attributes 882

Syntax Rules 882

CHAPTER 31 Optimizing Exchange 2000 *(continued)*

Modifying the Active Directory Schema 882
Using the Active Directory Schema Snap-in 882
Scripting 882
Deactivating Schema Components 883
Implications of Modifying the Schema 884
Optimizing Exchange 2000 Server 885
Scaling Front-End and Back-End Servers 886
Creating Multiple Virtual Servers 886
Full-Text Indexing 887
Searches 888
Building the Index 889
Gather Files 889
Scheduled Updates 890
Supporting Clients 890
Best Practices for Full-Text Indexing 890
Configuring DNS for a Unified Namespace 892
Accessing Active Directory Data 893
Outlook 2000 893
Compatibility with Earlier Versions of Outlook 894
Server Directory Access 894
Making Bulk Changes to Active Directory 895
Storing Data in Active Directory 895
Data Partitions in Active Directory 895
Global Address List 897
Selecting Attributes to Replicate to the Global Catalog 897
Preparing for Administration 898
Configuring an Administrator Computer 899
Administering Exchange 2000 and Exchange Server 5.5 899

CHAPTER 31 Optimizing Exchange 2000 *(continued)*

Tuning Exchange 2000 Performance 900
DSAccess Settings 900
DSAccess Cache 900
DSAccess Configuration 903
Relocation of Database, Log, and STM files 904
Optimizing Message Transport 904
Overview of Message Routing in Exchange 2000 905
Routing Groups 905
Connectors 906
Link State 906
Message Routing and Group Expansion 906
When to Use Multiple Routing Groups 907
Routing Group Topology 907
Multiple Paths Available Between Routing Groups 907
Message Traffic Analysis Can Help Define Boundaries 907
Additional Considerations 908
Connecting Routing Groups 908
Planning Routing Group Boundaries 908
Routing Group Deployment Example 909
Routing Group Connector 911
SMTP Connector 912
Link State Table 913
Link State Information 913
Routing Group Masters 914
Link State Algorithm 914
Link State Table Maintenance 914

CHAPTER 31 Optimizing Exchange 2000 *(continued)*

Message Routing 915
Routing Messages Within the Same Server 916
Routing Messages Within the Same Routing Group 916
Routing Messages to a Server in a Different Routing Group 917
Determining a Route Through Multiple Routing Groups 917
Routing Messages Outside an Exchange 2000 Organization 919
Rerouting Mail 919
Retries 920
Multiple Public Folder Trees 921
Managing Public Folders 922
Public Folders in Active Directory 922
Public Folder Replicas 923
Connecting to a Public Folder Replica 923
Public Folder Referrals 924

CHAPTER 32 Real-Time Collaboration 925

Instant Messaging 925
Instant Messaging Components 926
Instant Messaging Server Components 926
Instant Messaging Client Components 931
Client Operations 931
Instant Messaging Addressing 932
Windows 2000 Dependencies 934
Configuring Instant Messaging 934
Configuring DNS 936
Configuring Server Components 938
Instant Messaging Server Limits 939

CHAPTER 32 Real-Time Collaboration *(continued)*

Instant Messaging Deployment Scenarios . . . 940
Small Business Deployment . . . 940
Standard Deployment . . . 941
Enterprise Deployment . . . 941
Multiple E-mail Domain Deployment . . . 942
ISP Deployment . . . 943
Chat Service . . . 944
Server Configuration . . . 944
Scalability . . . 945
Client Connections . . . 945
Exchange Conferencing Server . . . 945
CHAPTER 33 Troubleshooting . . . 947
Installation and Setup Problems . . . 948
Setup Fails: "Setup Was Unable to Bind to the Exchange Server" Error Message . . . 948
Solution . . . 949
Setup Fails When You Install Exchange in a Child Domain . . . 949
Solution . . . 949
Setup Fails: "Error 0xC103798A" Error Message Displays . . . 949
Solution . . . 949
Setup Fails While Trying to Join Existing Exchange 5.5 Site . . . 949
Solution . . . 950
Setup Fails While Joining an Exchange Server 5.5 Site on a Windows 2000 Domain Controller . . . 950
Solution . . . 950
False Alerts, Server Reboots, Services Restart That Don't Exist When Upgrading from Exchange 5.5 . . . 950
Solution . . . 951
Setup Fails When Adding a New Server to a Site . . . 951
Solution . . . 951

CHAPTER 33 Troubleshooting *(continued)*

Public Folders Fail to Replicate . . . 952
Solution . . . 952
URL Not Available During Initial OnCreate Event . . . 952
Solution . . . 952
Problems With Active Directory and Active Directory Connector . . . 953
Cannot Modify Active Directory Objects . . . 953
Solution . . . 953
Exchange 2000 Options Unavailable in Active Directory Users and Computers . . . 953
Solution One . . . 953
Solution Two . . . 953
Hidden Objects in Exchange 5.5 are Visible in Exchange 2000 . . . 954
Solution . . . 954
ADC Creates a Configuration Connection Agreement . . . 954
Solution . . . 954
Active Directory Connector Setup Doesn't Update Schema . . . 955
Solution . . . 955
Web Storage System Problems . . . 955
ExIFS Doesn't Start After Stopping Manually . . . 955
Solution . . . 956
Remote File Operation Against an IFS Share Fails With Error 0xEFAD2521 . . . 956
Solution . . . 956
"System Cannot Find the Path Specified" or "x:\mailbox is Not accessible" Error . . . 957
Solution . . . 957
Last Access Time Is Not Updated On Files . . . 957
Solution . . . 957
Cannot Mount an Additional Web Storage System . . . 957
Solution One . . . 957
Solution Two . . . 958
Solution Three . . . 958

CHAPTER 33 Troubleshooting *(continued)*

Exchange 5.5 to Exchange 2000 Mailbox Move Fails . . . 958
Solution . . . 959
Client Displays Incorrect Public Folders After Modifying Public Database . . . 959
Solution . . . 959
Mount or Dismount All Databases Option Missing in a Storage Group . . . 959
Solution . . . 960
Connectivity Problems . . . 960
Outlook Client Messages Are Lost . . . 960
Solution . . . 960
Ics.dat File is Corrupted or Lost . . . 963
Solution . . . 963
Site Replication Service Does Not Run in Native Exchange 2000 Environment . . . 963
Solution . . . 963
Instant Messaging Logon Fails in Exchange 2000 but Works in Windows 2000 . . . 964
Solution . . . 964
Instant Messaging Contacts Do Not Display in an Online Conference . . . 964
Solution . . . 964
Instant Messaging Client Disconnects When Changing Status . . . 965
Solution . . . 965
"Conference Not Found" Error Message Displays Using Word 2000 as E-mail Editor . . . 965
Solution . . . 965
Deleted Newsgroup Remains in Exchange 2000 Public Store . . . 965
Solution . . . 966
"Deleted Server" Appears in Replication Monitor . . . 966
Solution . . . 966
Configuring SMTP Connector to Run on Non-Standard Port to Enhance Security Does Not Work . . . 966
Solution . . . 966

CHAPTER 33 Troubleshooting *(continued)*

SMTP Error: "Recipient Could Not Be Reached 967
Solution 967
Messages Not Flowing Between Servers in Different Routing Groups 967
Solution 968
SMTP Protocol Error "454 Client Must Re Authenticated" 968
Solution One 968
Solution Two 968
Cannot Gain Access To Additional HTTP Virtual Server on Cluster 968
Solution 968
Windows 2000 Server Tools Problems 969
Cannot Access Registry 969
Solution 969
Backup and Restore Problems 970
Cannot Create Identical Mailboxes in a Site or Organization 970
Solution 970
Event ID 1018,1019, and 1022: Database is Damaged 970
Solution 970
Storage Group Fails to Mount with -1216 (JET_errAttachedDatabaseMismatch) or Oxfffffb40 Error 970
Solution One 971
Solution Two 971
Performance Problems 971
Memory Leak in LDAP Service When Using Migration Wizard 972
Solution 972
APPENDIX B Ports and Protocols 973

CHAPTER 26

Exchange 2000 Architecture

Microsoft Exchange 2000 Server architecture utilizes the Microsoft Windows 2000 operating system. in particular the Active Directory directory service and Simple Mail Transfer Protocol (SMTP) native transport, to provide a scalable and efficient messaging system. Exchange 2000 provides event support for protocol, transport, and Microsoft Web Storage System to allow for greater extensibility. The Exchange 2000 client-side architecture efficiently manages messaging and publishing information, and you can use real-time collaboration architecture to transfer dynamic data and information on demand.

In This Chapter

Server Architecture

Message Flow

Outlook Web Access Architecture

Exchange 2000 and Previous Versions of Exchange

Server Architecture

Exchange 2000 Server architecture provides users with a scalable messaging and data retrieval and sharing system. It allows for greater versatility in network design because it integrates with Windows 2000 Active Directory, Microsoft Internet Information Service (IIS), and Web Storage System. Protocols such as Network News Transport Protocol (NNTP), Post Office Protocol version 3 (POP3), and Internet Message Access Protocol version 4 (IMAP4) run as part of the IIS process, allowing Exchange 2000 architects to dedicate servers to specific tasks, such as managing the databases or processing client requests. The use of multiple databases and protocols expands the server's capabilities to reduce network traffic, expedite and broaden data retrieval and usage, and provide scalability, security, and rapid recovery and failover.

Exchange 2000 supports multimedia formats, making information storage and retrieval fast and efficient. It stores native content, including Multipurpose Internet Mail Extensions (MIME) content, directly in an .stm file. Also, with the installable file system (IFS), you can use Exchange as a file repository for any application. IFS makes it possible to map Exchange folders and mailboxes as shared network drives. Finally, Exchange supports multiple public folder hierarchies;

each hierarchy (or tree) is stored in a public folder store. The default server installation includes one public folder store that contains one public folder hierarchy. Figure 26.1 shows a graphical representation of the Exchange 2000 architecture, including the Extensible Storage Engine, Web Storage System, Exchange Interprocess Communication Layer (ExIPC), and IIS, and their underlying components.

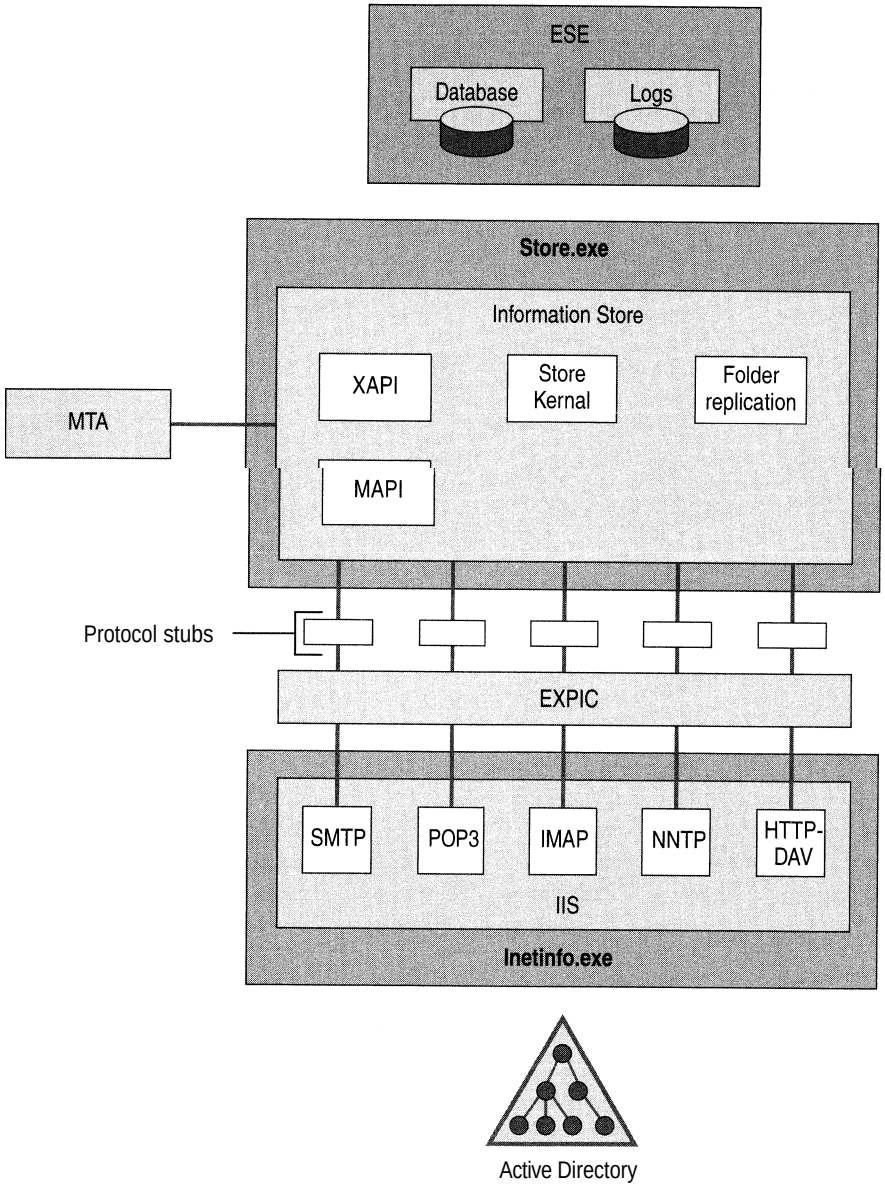


Figure 26.1 Exchange 2000 architecture

IIS Process

The protocols in Exchange 2000 communicate with Web Storage System when they run as part of the IIS process. The IIS process must communicate efficiently and rapidly with Web Storage System to provide clients with immediate and reliable access to data in the store. The ExIPC serves this function.

All Exchange 2000 protocols are hosted within the IIS process. Exchange expands the SMTP service, enhancing the basic delivery functions of the protocol without compromising its compatibility with other messaging systems. Exchange gives you more control over the routing and delivery of messages, and provides secure access and channels for managing the SMTP service.

Web Storage System

Exchange 2000 sub-systems, such as protocol, storage, and directory features, can be placed on different servers to improve scalability. For example. you can configure Exchange 2000 topology to consist of a bank of front-end and back-end servers, enabling users to connect to virtual Internet Protocol (IP)-addressable front-end servers while storing messages and collaboration data on separate back-end servers.

You can build large public e-mail systems by separating components onto different servers. This increases reliability because a failure on one server will not affect the other components. For example, you can host a set of protocol servers separately without Redundant Array of Inexpensive Disks (RAID) controllers to reduce cost and prevent protocol failures from affecting Web Storage System or Active Directory. Because Exchange 2000 doesn't require that you host directory services on each computer. you have more flexibility in your deployment.

> **Note** MAPI clients, such as Microsoft Outlook 97, Microsoft Outlook 98, and Microsoft Outlook 2000 clients working in a "Corporate Workgroup" mode, cannot use a front-end server; they must always use the back-end server.

Exchange Interprocess Communication Layer

Although decoupling the protocols from Web Storage System increases the reliability, flexibility, and scalability of Exchange 2000, you must have a very fast method for the protocols to exchange data with Web Storage System. To facilitate the rapid transfer of information between the IIS process and Web Storage System, Exchange 2000 has a queuing layer called the Exchange Interprocess Communication Layer (ExIPC) that allows the IIS and Web Storage System processes to quickly exchange data.

Figure 26.2 shows a graphical representation of the ExIPC, including circular queues, protocol services in IIS (such as smtpsvc), and the corresponding ExIPC interface, or "protocol stub," in Web Storage System (such as exsmtp).

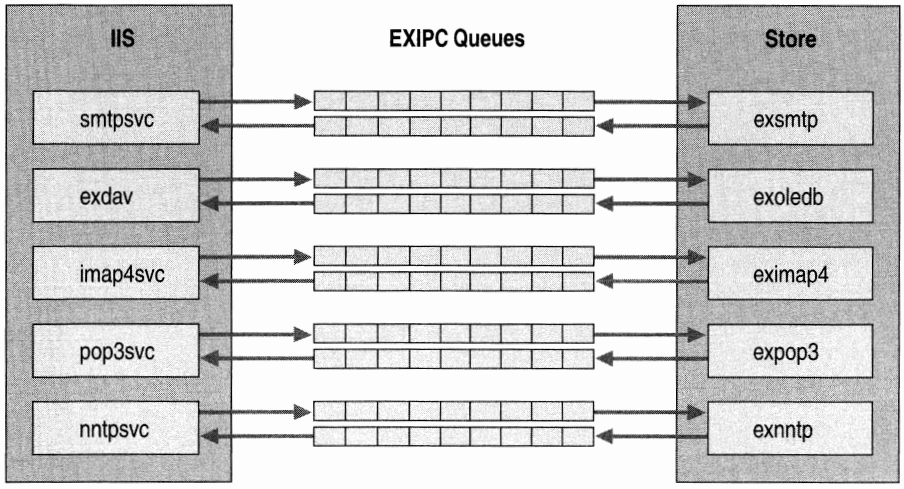


Figure 26.2 Exchange Interprocess Communications Layer (ExIPC) architecture

ExIPC is a high-performance interprocess communication facility. ExIPC uses shared memory from a shared memory heap to communicate between processes. ExIPC consists of a protocol dynamic-link library (DLL) that implements a binding facility, a shared memory heap, and a pair of queues based on shared memory.

> **Note** A *heap* is a portion of memory reserved for a program to temporarily store data structures whose existence or size is unknown until the program runs. The program can request free memory from the heap to hold such elements, use it as necessary, and later free this memory.

One of the key advantage\ of ExIPC is that, unlike most lightweight remote procedure calls (LRPCs), it functions asynchronously. This allows Exchange to allocate memory immediately after one portion of a process completes, which results in corresponding improvements in performance.

ExIPC Binding Facility

The ExIPC binding facility allows you to create and connect an arbitrary number of queues between two processes, such as the IIS and Web Storage System. This binding facility includes the Central Queue Manager that keeps track of the queues and processes with which a process is communicating. The facility is also used for unbinding and queue clean up in the event of a catastrophic failure on the other process.

Each protocol queue is circular and of a fixed size. During interprocess communications, data is stored in the shared memory heap, referenced by a data handle. The data handle is enqueued and dequeued, and the IIS or the store then references the portion of shared memory from the handle.

ExIPC Protocol Stubs

Each protocol has an ExIPC interface in Web Storage System. For example, the ExIPC protocol stub for POP3 is expop.dll. This component passes parameters (for example, a pointer to a message or an action) from Web Storage System to the ExIPC interface (drviis.dll) in the IIS process.

Sample of ExIPC Process

Figure 26.3 is a graphical representation of the ExIPC process for when a client checks for a message in Web Storage System.

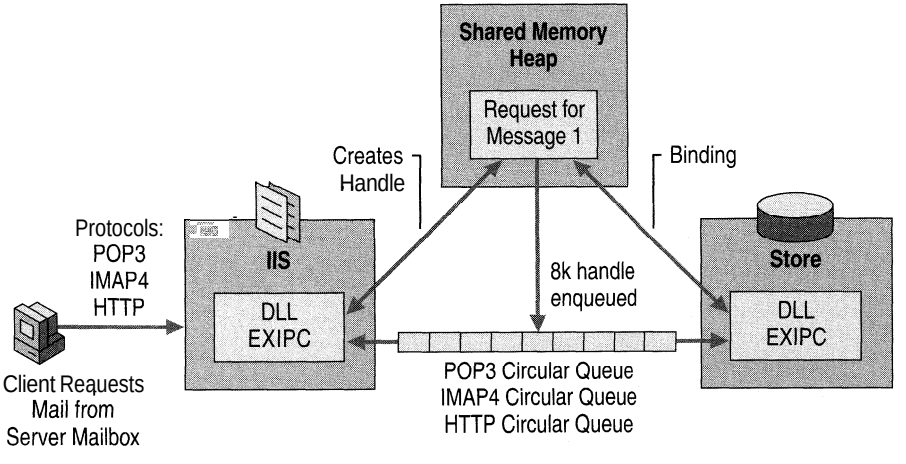


Figure 26.3 Exchange Interprocess Communications Layer process

The following steps illustrate the interprocess communication steps that ExIPC goes through when a client checks a mailbox on the server using a POP3 protocol application (such as Microsoft Outlook).

1. The client logs on to the server and gives the command to check e-mail.
2. A Request Mail Message #1 command is created on the IIS side.
3. IIS allocates an amount of shared memory from the shared memory heap for the request. A corresponding handle is assigned to that portion of the shared memory. The handle, which functions as a placeholder or pointer to a referenced portion of memory, is then placed into the circular memory queue (enqueued) in the direction of the Store.
4. On the Store side, the ExIPC DLL for POP3 checks for incoming POP3 requests. The DLL receives the Request Mail Message and removes the handle from the circular memory queue. The Store-side POP3 stub references the handle to the data in the shared memory heap.
5. If there are no failures or performance problems on the Store side, the ExIPC process is complete and the data is successfully communicated from the IIS to the Store. If a queue is full or the store has stopped, an error message is returned.
6. A response (the mail message) is generated on the Store side. The Store allocates the appropriate amount of shared memory for the response from the shared memory heap. A corresponding handle is assigned to that shared memory. The handle is then enqueued in the direction of the IIS.
7. The IIS removes the handle from the circular queue, references the shared memory, and binds them together.
8. If there are no failures or performance problems on the IIS side, the response is complete and the data is successfully communicated from the Store to IIS.

Supported Internet Protocols

The protocols that Microsoft Exchange 2000 uses to access Web Storage System, such as SMTP, IMAP4, NNTP, POP3, and HTTP, have been redesigned in Exchange 2000 and are now integrated with IIS in Windows 2000. This redesign improves scalability and adds administrative flexibility.

The protocols are administered using Exchange System Manager, rather than Internet Services Manager. Configuration is saved to Active Directory and then replicated to the IIS metabase, on the appropriate Exchange 2000 server, by the Exchange System Attendant.

Exchange Server 5.5 and Outlook already provide native support for Internet messaging and collaboration standards, including Hypertext Transfer Protocol (HTTP), SMTP, POP3, and IMAP. When you install Windows 2000, core Internet protocol stacks, such as SMTP and NNTP, are configured as part of the operating system. Windows 2000 uses these protocols for operations such as directory replication and basic message creation.

When you install Exchange 2000, it extends the core protocol stacks in Windows 2000 with additional command verbs and advanced routing components to provide all of the features required

for an enterprise-class messaging and collaborative system. The following sections describe how Exchange 2000 enhances the core Internet protocol stacks.

Simple Mail Transfer Protocol

The Simple Mail Transfer Protocol (SMTP) service processes incoming traffic from SMTP clients such as Microsoft Outlook Express and other SMTP hosts. It sends outbound SMTP traffic using configured SMTP Connectors and Routing Group Connectors.

The SMTP service is a key component in Exchange 2000 because it is the primary protocol for communicating with other computers running Exchange 2000, replacing the remote procedure call (RPC). It is also a key protocol for transferring e-mail over the Internet and providing interoperability with other e-mail systems.

> **Note** In previous versions of Exchange, the Internet Mail Service provided the SMTP engine features. The Windows 2000 SMTP service, which is a component of IIS, now provides these functions.

When you use the SMTP service with Exchange, the following changes to the SMTP service will occur:

- Messages destined for a local domain are now transferred by default to Web Storage System instead of a drop folder.
- Exchange System Manager configures the SMTP service. Exchange System Attendant applies this configuration, which is stored in Active Directory, to IIS.
- Exchange System Manager uses an SMTP service application programming interface (API) to display the status of mail messages queued for delivery.

SMTP Process

Most SMTP processing occurs in IIS rather than in the Internet Mail Service.

1. The SMTP engine component receives SMTP requests on TCP port 25 and creates an **ImailMsg** object for incoming messages.
2. The **ImailMsg** object is passed to the advanced queuing component, which places the message in the PreCatQueue and relays the message to the message categorizer component. The message is returned after its destination has been determined.
3. The **ImailMsg** object is placed in the PreRoutingQueue when the route for the message is determined. The Advanced Queuing Component then queues the message for delivery, taking into account next-hop routing information supplied by the routing engine.
4. The Advanced Queuing Component delivers the message either to the local Web Storage System or causes the SMTP service to send the message outbound using SMTP. The message is delivered to Web Storage System in the event that it is destined for a local recipient, public folder, or the Message Transfer Agent for transfer to other Exchange 5.5 or earlier servers through RPC, or X.400. The protocol portion of the SMTP service processes regular outbound messages for other Exchange 2000 servers through the SMTP protocol.

Figure 26.4 shows a graphical representation of the Exchange 2000 SMTP process.

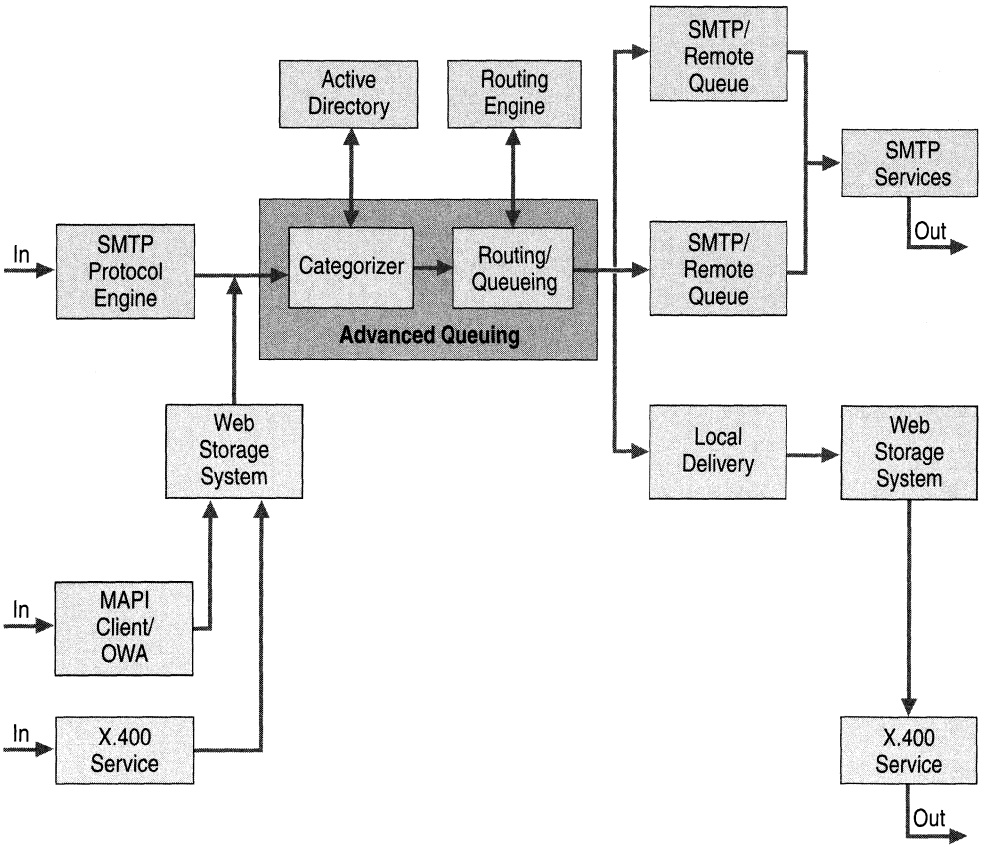


Figure 26.4 SMTP process

Windows 2000 Server includes a native SMTP component designed to be a basic implementation of the protocol. Windows 2000 and other products can use this transport to perform operations such as directory replication. Because SMTP is part of IIS 5.0, it installs with Windows 2000 and it is extended in Exchange 2000. It runs as part of the Inetinfo.exe process. It supports basic distribution list expansion features. It cannot act as a ListServer, and it cannot route messages based on topology link status.

SMTP Directories

The SMTP root directory is root:\Inetpub\Mailroot.

Within the root directory are four folders:

- **Drop** Stores incoming messages from all domains hosted on the computer. You can assign any directory to be the Drop directory, except one that has already been designated as the Pickup directory. This feature is not available on an Exchange 2000-based server.

- **Pickup** Processes outgoing messages that are created as text files and copied to the directory. When a properly formatted RFC 822 message is copied to the Pickup directory, the SMTP service collects it and initiates delivery.
- **Queue** Holds messages waiting for delivery when they cannot be delivered because the connection is busy or down.
- **BadMail** Stores undeliverable messages that cannot be returned to sender.

SMTP Transport Architecture

Windows 2000 and other products can make use of native SMTP component transport to perform operations such as directory replication.

Figure 26.5 is a graphical representation of the Exchange 2000 SMTP transport architecture.

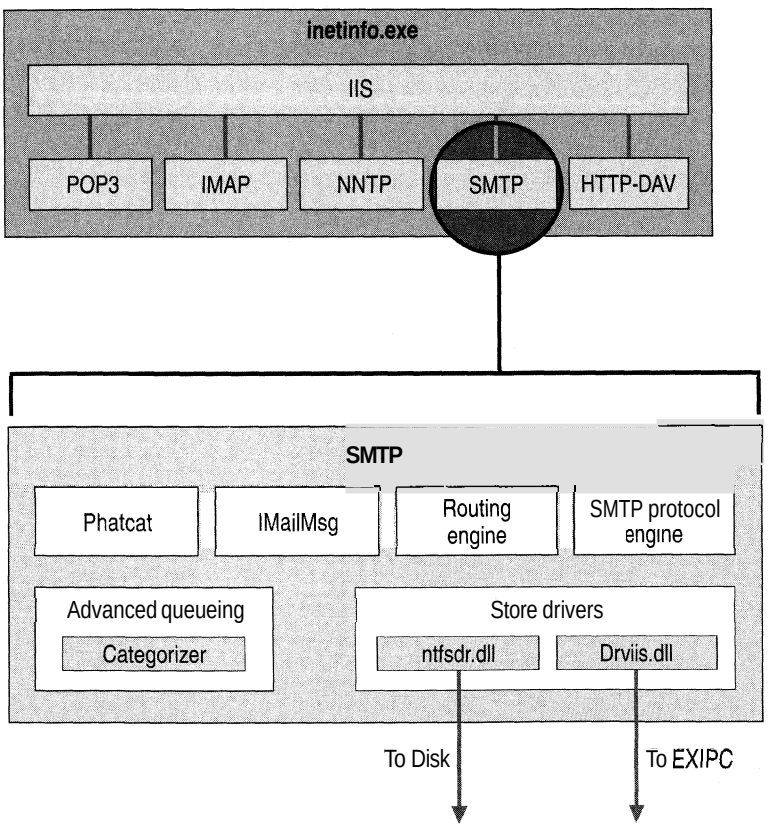


Figure 26.5 Exchange 2000 SMTP transport architecture

SMTP Extensions

When Exchange 2000 is installed, it extends the features of the SMTP service and adds new components through the use of transport and protocol event sinks. Some of the key extensions that Exchange 2000 adds to SMTP are:

- Command verbs that support link state information
- An advanced queuing engine
- An enhanced messaging categorization agent

Exchange 2000 enhancements to the SMTP service are described in the following sections.

SMTP Protocol Engine (Smtpsvc.dll)

The SMTP Protocol Engine (Smtpsvc.dll) initiates outbound connections and answers inbound connections. The DLL (Smtpsvc.dll) forms the core of the SMTP service and also handles the issuing and processing of commands.

ImailMsg

ImailMsg obtains a file handle to the message from the \Exchsrvr\Vs1\Mailroot\Queue directory. This component also preserves P1 information as part of the message in a separate data stream.

When Exchange 2000 is installed, it adds another store driver (Drviis.dll). As a result, there is a pointer to the message in Web Storage System when the message goes to an Exchange 2000 recipient, or a file handle when the message goes to a drop directory.

Store Drivers

The store driver saves the message to the hard disk drive using Ntfsdrv.dll. The messages are saved to the queue directory in the \Exchsrvr\Vs1\Mailroot directory. If the message is addressed to a local domain, the message is then moved to the specified drop directory, which is \Exchsrvr\Vs1\Mailroot Drop by default.

When Exchange 2000 is installed, it adds another store driver (Drviis.dll) to communicate with Web Storage System by using ExIPC. After adding a new local domain, the system is capable of delivering messages to either Web Storage System (using Drviis.dll) or the file system (using Ntfsdrv.dll).

Advanced Queuing Engine

The advanced queuing engine defines and manages queues for message delivery, such as domain queues and link objects that you can query for transport information. Once the advanced queuing engine receives an SMTP message, it relays the message to the categorizer. The categorizer returns the message after it determines the destination. The advanced queuing engine then queues the message for eventual delivery, taking into account remote routing information from the routing engine. Lastly, it delivers the message to either the local Web Storage System driver or the local SMTP stack.

Message Categorizer

The message categorizer is a plug-in to the advanced queuing engine. It is essentially a collection of event sinks that perform advanced address resolution on every message that travels through the advanced queuing engine. These messages might be destined for the local Web Storage System, a remote host through the message transfer agent (MTA), or a remote host through SMTP. The message categorizer also handles the expansion of distribution lists. The message categorizer is turned off by default in Windows 2000; Exchange 2000 activates it and enhances it with a set of categorizer event sinks, known as Phatcat, in Phatcat.dll.

Routing Engine

The routing engine adds link-state routing capabilities to queue architecture by providing accurate next-hop information to the advanced queuing engine.

Pipelining

Pipelining (RFC 2197) is a Windows 2000 SMTP service feature that allows multiple commands to stream from the SMTP client to the server. Pipelining enables the client to stream multiple commands to the server without waiting for server response, reducing the number of turnarounds in a client/server "conversation," thus increasing data transfer speed.

Table 26.1 shows an example of what happens in applications where Pipelining does not exist. The server can only process one command before responding. Each time that the client must wait for a server response is called a turnaround. In this example, there are nine turnarounds.

Table 26.1 When Pipelining is not enabled

Action	Turnarounds
Server: <wait for open connection>	
Client: <open connection to server> S: 220 domain.com SMTP service ready	1
C: HELO nwtraders.microsoft.com S: 250 nwtraders.microsoft.com	2
C: MAIL FROM:<user@nwtraders.microsoft.com S: 250 sender<user@nwtraders.mircosoft.com> OK	3
C: RCPT TO:<user2@domain.com> S: 250 sender <user2@domain.com>OK	4
C: RCPT TO:<user3@domain.com> S: 250 sender <user3@domain.com>OK	5
C: RCPT TO:<user4@domain.com> S: 250 sender <user4@domain.com>OK	6
C: DATA S: 354 enter mail, end with line containing only "."	7
C: . S: 250 message sent	8
C: QUIT S: 211 goodbye	9

Table 26.2 shows an example of when pipelining is enabled. Because the server can process multiple client commands without requiring a response, the turnarounds are reduced from 9 to 4.

Table 26.2 When Pipelining is enabled

Action	Turnaround
Server: <wait for open connection>	
Client: <open connection to server>	1
S: 220 domain.com SMTP service ready	
C: HELO nwtraders.microsoft.com	2
S: 250 nwtraders.microsoft.com	
S: 250 PIPELINING	
C: MAIL FROM:<user@nwtraders.microsoft.comC: RCPT TO:<user2@domain.com>	
C: RCPT TO:<user3@domain.com>	
C: RCPT TO:<user4@domain.com>	
C: DATA	3
S: 250 sender<user@nwtraders.mircosoft.com> OK	
S: 250 sender <user2@domain.com>OK	
S: 250 sender <user3@domain.com>OK	
S: 250 sender <user4@domain.com>OK	
S: 354 enter mail, end with line containing only "."	
C: .	
C: QUIT	4
S: 250 message sent	
S: 221 goodbye	

Chunking

Chunking (RFC 1830) is a Windows 2000 SMTP service feature that enhances the way both large and binary messages are sent via SMTP. Chunking uses an SMTP verb called BDAT as an alternative to the DATA command (RFC 821). In DATA command arguments, the end of a data stream is indicated as line feed, full stop, or carriage return–linefeed. When the host computer receives data, it must scan for the end of the data representation, creating excessive network traffic. The BDAT command tells the host computer how many octets of data are contained in the message, eliminating the need to scan the data and increasing data transfer efficiency. In the following example, BDAT 70 indicates that the length of the binary data packet is 70 octets. LAST indicates which data packet is the last one.

```
Client: <wait for connection on TCP port 25>
Server: <open connection to server>
C: 200 nwtraders.microsoft.comSMTP service ready
S: EHLO nwtraders.microsoft.com
```

```
C: 250-nwtraders.Microsoft.com
C: 250 CHUNKING
S: MAIL FROM: <user1@domain.com>
C: 250 <uer1@domain.com>...Sender ok
S: RCPT TO: <user2@domain.com>
C: 250 <user2@domain.com>...Recipient ok
S: BDAT 70 LAST
S: To: user2@domain.com<CR><LF>
S: From: user1@domain.com<CR><LF>
C: 250 Message OK, 70 octets received
S: QUIT
C: 221 goodbye
```

POP3 and IMAP4

As with the other protocols, POP3 and IMAP4 are integrated with IIS. You use the Exchange System Manager to configure each virtual server with the settings stored in Active Directory. The Exchange System Attendant applies this configuration to IIS.

Exchange 2000 enhances POP3 and IMAP4 services with the following features:

- **Virtual sewers** You can configure servers with separate names, authentication, or message formatting.
- **Front-end/back-end sewers** You can use a single namespace for multiple servers. Clients connect to the front-end server, which looks up the user's mailbox in the directory and then proxies the traffic to the corresponding back-end server. The front-end server also provides IMAP4 clients access to all public folders, even those that do not exist on the primary public server.
- **IMAP4 support for RFC 2359** **RFC** 2359 describes how to reduce server communication for copied and appended messages. Using this protocol, you can gain access to public folders, grant Delegate permissions to another user's mailbox, allow anonymous access to specific IMAP4 account names, and store messages submitted by an Internet client in native MIME format.

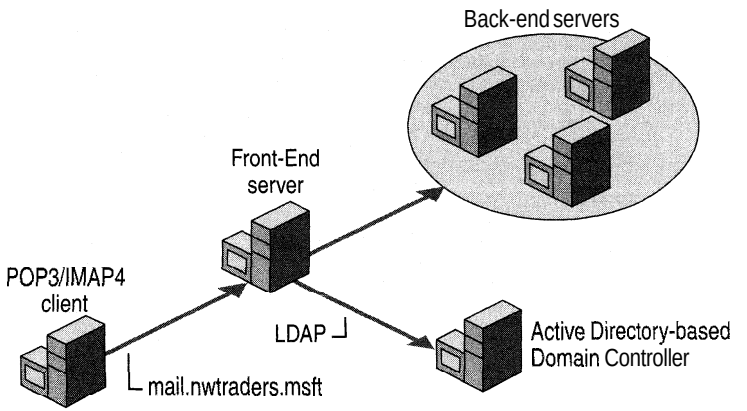


Figure 26.6 Exchange 2000 POP3/IMAP4 front-end/back-end configuration

NNTP Support

The NNTP service in Windows 2000 is designed to support a stand-alone newsgroup server that can create group discussions easily. When Exchange 2000 is installed, the NNTP service is enhanced with the ability to interface with other news servers through news feeds. The NNTP service can communicate with external NNTP servers to make popular USENET groups available to users.

The standard storage location for the NNTP service is one or more directories in the file system. With Exchange 2000, the NNTP service can also store newsgroups in the public folders of any available public folder tree. The default location for newsgroups is the Internet Newsgroups folder. The NNTP service uses virtual directories to reference these locations.

You can arrange multiple news servers in a master server/subordinate server layout. This enables clients to connect to a large group of servers and still maintain accurate views of newsgroup content. Having a bank or group of servers provides additional scalability for a large number of clients and provides fault tolerance if a subordinate server goes offline.

The Exchange 2000 implementation of NNTP provides the following additional features for this protocol:

- Content indexing provides search features for public folders
- Full news feeds accepted independent of back-end storage
- MAPI or NNTP clients can read or post to newsgroups supported by Web Storage System

Figure 26.7 shows the NNTP architecture.

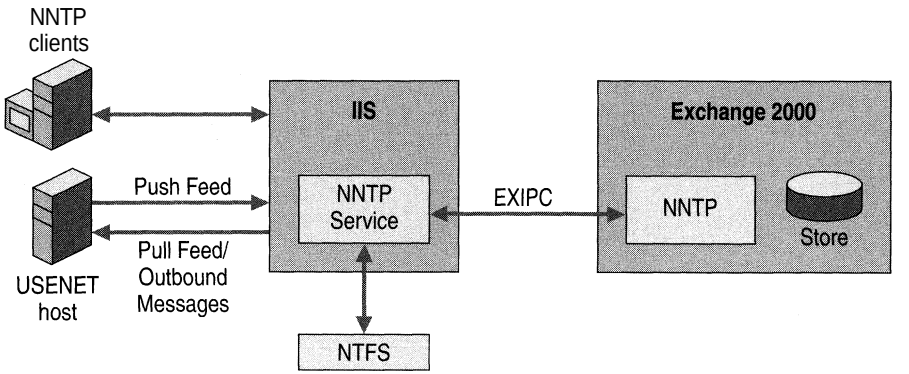


Figure 26.7 NNTP architecture

Web Storage System

Web Storage System in Exchange 2000 makes access to data fast and easy. Web Storage System brings the Web, the file system, and the collaboration sewer into a single place for storing information and deploying applications. It uses features in previous versions of Exchange such as transaction logging, single instance storage, rollback recovery, and online maintenance and integrates them with the Web to provide users with more integrated information management options.

The Exchange 2000 architecture gives you the opportunity to use very large databases for large-scale enterprises and the option of splitting a single logical database into separate physical databases to increase overall system reliability, enable faster backup, and quicken the recovery process in the event of hardware failure. For example, if the hardware running a part of the overall e-mail database fails, only that database is affected during restore while others continue to serve their e-mail users, even though the database is administered as a single unit.

Web Storage System includes the following features:

- A database managed as a hierarchal file system of folders and items
- Support for native Web content
- Accessibility to all information using common URLs
- Consistent access from Outlook or any e-mail client, and Windows Explorer or any Web browser
- Built-in content indexing and search capabilities
- Support for several protocols and APIs for remote client access

Microsoft Search Service

Web Storage System uses Microsoft Search Service, which can index Microsoft Word, Microsoft Excel, and Microsoft Office documents, HTML files, text files, and the text portions of .ppt files. You can also search with Ifilters. which enable indexing of .pdf files. Searches are transparent to the user, and Exchange provides an immediate list of hits, complete with attachments and pertinent text passages.

All searches must pass through Exchange 2000, which is responsible for all security enforcement. The Exchange Query Processor processes query splitting and result merging.

Storage Location for Documents and Applications

The main benefit of Web Storage System is that it provides a single place where you can store and manage documents, collaborate as a team, and run powerful business applications. For example, as more applications move to the Web, back-end server functions increase. Before Web Storage System existed, you needed to deploy and integrate multiple services to meet the needs of both developers and end-users. If you wanted to implement a simple tracking application, for example, you needed the following services:

- **File System** This was required to allow clients such as Microsoft Office to read and write documents and store streaming data such as audio and video. It had to include data models that supported hierarchal collections such as folders and heterogeneous collections such as multiple item types that could be placed in folders.
- **Database Services** These services were required for queries beyond what was possible on files in the file system. Database services were also required for applications to present a consistent view to users when updates involved more than one item in Web Storage System.
- **Collaboration Services** A tracking application that requires messaging, contact, and calendaring support.
- **Content Indexing** A content indexing application was required; otherwise, you would need to implement Site Server to have content indexing features.

Web Storage System delivers all of these services in one integrated package, making it easier to find, use, and share information. Web Storage System allows application designers to expand their use of Exchange 2000, Office 2000, and Microsoft BackOffice as a platform for information management.

Support for Multiple Databases

Exchange 2000 provides support for multiple databases and storage groups on the same server. An Exchange storage group is part of the Extensible Storage Engine (ESE). Exchange 2000 allows up to five multiple databases per storage group, with four storage groups possible per server. Each database must be homed inside a storage group. In turn, each storage group is contained in a virtual cerver. On a single node server with no clustering, there is only one virtual server running.

However, if you have two nodes inside a cluster, you can create more virtual servers and independently fail them over if a node fails.

Note When creating virtual servers in a cluster, the total number of virtual servers should be (n-1), where n=the number of nodes. This leaves one free node for fail over.

Although each instance of a database runs under the same Web Storage System process, you can mount or dismount individual databases dynamically, in other words run only one process per server. The advantage is that you can restore an individual database from backup while other database instances service client requests. With Exchange 2000's multiple database architecture, each database is checked when the Web Storage System process begins; if one or more databases is corrupt or unavailable, the Web Storage System process can continue, allowing you to troubleshoot and restore the affected files while the system remains up and running.

Multiple databases allow you to increase the number of simultaneous users maintained on a single server and reduce the impact of a failure. Multiple databases allow you to strategically plan your organization's data storage by classifying various kinds of data for each database and assigning separate databases for your organization's most important users. This allows you to easily and quickly recover vital data.

The multiple database architecture allows for front-end/back-end server configuration, splitting database storage engines (back-end) on separate servers from the protocols (front-end). This gives you the ability to use a unified namespace, and it increases security and scalability.

Note MAPI clients cannot access front-end servers and simultaneously use back-end servers.

Content Storage

When a computer running Exchange 2000 receives an Internet message, the native MIME content is left intact and the server stores this message in the native content store storage file. If a MAPI client, such as Outlook, tries to read data that resides in the native content store, the RTF/HTML converter converts the MIME content to a set of MAPI content on demand and passes the data to the MAPI client.

Each Exchange 2000 database in Web Storage System supports three different file formats: log files, .stm files, and .edb files. An .stm file and an .edb file form a unit, which represents a single message that is divided into its native content (the .stm portion) and its space usage information and checksums (the .edb portion).

Native Content Storage Database Format

The native content store uses 4-kilobyte (KB) pages and can allocate data between non-contiguous areas, in much the same way as a rich-text store. The native content store performs online defragmenting and compacting, as does the rich-text store.

Rich-text Store

The rich-text store is the same database used in previous versions of Exchange. The rich-text store stores messages from MAPI clients, such as Outlook, the same way they were stored in previous versions of Exchange. MAPI clients can then access these messages without conversion. However, if an Internet protocol-based client attempts to read a rich-text message, IMAIL will convert the message to the requested format, using on-demand content conversion.

On-Demand Content Conversion in Exchange 2000

Exchange 2000 performs content conversion only when it is absolutely necessary. This process is known as *deferred content conversion.* For example, when an Internet protocol message arrives, the data passes through the native content store, and the native MIME streams are left intact. When another Internet protocol-based client attempts to read the data, it streams directly out of the native content store without conversion.

The on-demand content conversion process occurs entirely within the memory of a computer running Exchange 2000. The data does not physically move and it is not copied between databases. The exception to this is if the data changes. For example, if a client, such as Outlook, reads a message in the native content store and IMAIL converts the data in memory, and the user saves changes to the message, it is then submitted to the rich-text store, not the native content store. The original message still exists in the native content store if other users have pointers to it.

Installable File System

The Exchange IFS permits normal network client redirectors to share folders and items. The benefit of this is that you can map a drive letter or NET USE to any public folder or even to your own mailbox. Because your local computer can assign a drive letter to these resources, standard applications such as Windows Explorer and Office 2000 can read and write data directly to the native content store.

Accessing Data Using IFS

One of the advantages of using IFS is that a client can access it without installing any special software other than the network redirector. The drive map on an Exchange 2000 server shows a number of devices, including the M: drive. This is the portal into the store for Win32 access. By default, drive M: is shared with a share name of BackOfficeStorage. Therefore clients can map a drive to \\Server\BackOfficeStorage.

Beneath the share, users see a folder named after the domain where the Exchange 2000 server is installed, and beneath this, there are two folders: MBX and Public Folders. Figure 26.8 shows the M: drive directory in Exchange, including the MBX and Public Folders folder in the right pane.

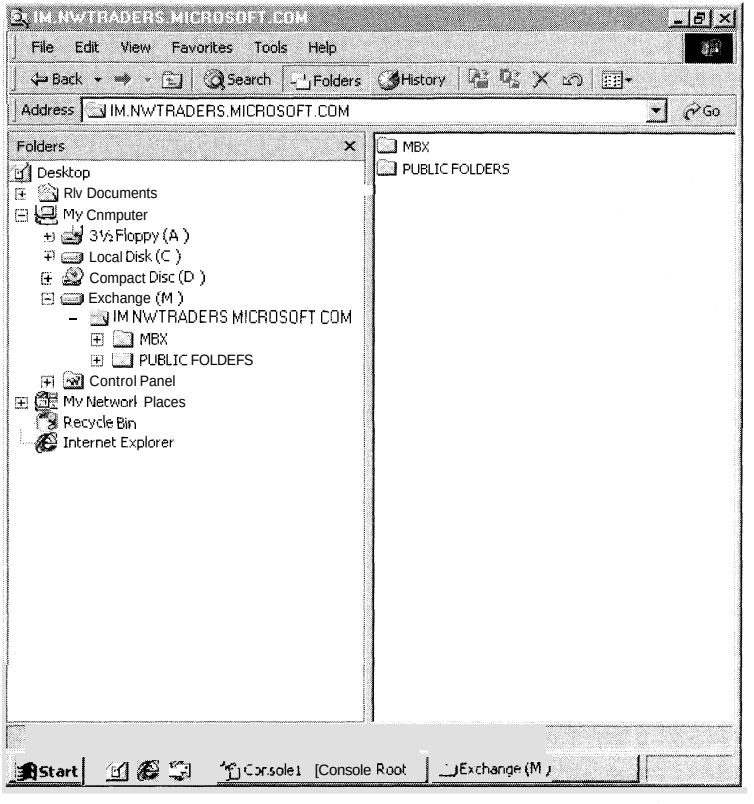


Figure 26.8 Accessing files through the Installable File System interface

The MBX Folder

Although the MBX folder is the root for all mailboxes on the server, the folder names are invisible. Win32 clients that have the appropriate access permissions can navigate to this part of the hierarchy using explicit folder names that are the same as mailbox alias names.

You can enumerate all objects in the Inbox with the DIR command and configure them with other file system operations, such as COPY and DELETE. All normal message objects are represented with .eml extensions and you can see the message stream itself using the TYPE command.

Public Folders Hierarchy

The public folders portion of the tree allows you to navigate the public folder hierarchy and access data within each public folder to which you have access.

ExIFS: Inbound Process

For inbound messages, the IIS creates a new file in ExIFS and writes to it. ExIFS then returns a list of pages that the message was written to. The list of pages passes to the Exchange Server process. The Extensible Storage Engine commits the pages by logging the information. The page checksums are stored in the EDB file.

ExIFS: Qutbound Process

For outbound messages, the Exchange server gets the list of pages referenced by a message and passes them to the IIS. which opens a file in ExIFS. The message is quickly transmitted. Page checksums are not verified during the transmission.

ExIFS: Space Allocation

The Extensible Storage Engine tracks which pages in the streaming file are committed. It reserves these pages and transfers them to ExIFS. The ExIFS allocates space for new messages from its reserved space. ExIFS also requests space from the Extensible Storage Engine when necessary.

Enhanced Active/Active Clustering

Exchange 2000 allows up to four nodes in a cluster, all of them running an instance of the Store.exe file in an active/active configuration. Because all of the nodes are consistently active, each node owns only a portion of the data load at one time. Therefore, if one of the nodes fails, there is no lag time required for another node to run Store.exe, and there are fewer storage groups to move. The storage groups on the failed node are immediately divided among the remaining nodes, which allows faster file swapping and less drain on resources.

Public Folders

A public folder stores shared user messages or information. Public folders can contain different types of information, from custom forms to Internet content stored in its native format. Public folders are a part of Web Storage System accessible by Exchange clients, Web browsers, and custom applications. You can control access to folders by setting permissions, and optimize their access and performance by replicating folders to other servers in your organization. Public folders contain many features that make Exchange a complete collaboration server.

Multiple Store Support

You can configure the public folder hierarchy into smaller logical units to reflect departments in an organization. In addition to the default public folder hierarchy, you can create other top-level folder hierarchies. Each folder hierarchy is represented by its own database in Web Storage System. Web Storage System supports multiple public stores running under a single process. Multiple stores increase scalability by allowing you to segment the folder data and distribute it across servers in your organization.

Web Access

Each public folder has an HTML page associated with it. When you create a folder, Exchange automatically creates a default HTML page. You can replace the automatically generated Web page with a custom page or you can change the URL to point to another Web site. You can gain access to all of the contents in a folder by accessing the folder container or by accessing a dynamically generated URL.

Replication

You can configure a public folder to be replicated on multiple servers. This is beneficial because you can distribute the user load among servers, distribute public folders across geographical areas, and back up public folder data. You can set up a replication schedule based on how often data in the public folder changes. You can set this schedule for all public folders or for a specific public folder.

Content Access on Alternate Sewers

When a client must use an alternate server to access public folder content, Exchange uses a routing group that is configured to refer the client to another server. This allows the client to access content when a specific server is not known. It also eliminates the need for a cost-based referral list. The connections between all servers in a routing group are likely to have the same bandwidth because routing groups are intended to serve as collections of well-connected servers.

Using Multiple Public Folder Trees

Exchange 2000 supports multiple public folder trees (also referred to as top-level hierarchies) for greater administrative control and flexibility of organization in public folders. For example, you can create a separate public folder tree for external users to keep content separate from the default public folder tree. You can also create an additional tree at a remote location for users there to access data that is relevant only to them.

When Exchange 2000 installs, it creates the default All Public Folders tree. This tree is available to all MAPI, IMAP4, NNTP, and HTTP Web clients. Public folder trees are only available to NNTP and Web clients, and not to clients such as Microsoft Outlook 2000 (unless viewed on a Web page hosted in Outlook 2000). You can use these non-MAPI accessible folders to communicate with browsers and applications that use HTTP to gain access to Web Storage System, such as Office 2000.

Each public folder tree stores its data in a single public folder store on each server. You can replicate folders in the tree to every server in the organization that has a public folder store associated with that public folder tree.

The multiple public folder feature can affect your public folder strategy. While the default public folder tree is created on every public folder server and its list of folders is always replicated, additional public folder trees only affect the servers on which they are configured. This means that you can create a set of departmental or local folders on only one or a subset of servers. You do not have to replicate these additional public folders to every public folder server. You can use additional public folder trees to minimize the overall size of the default public folder tree to simplify navigation and reduce the cost to replicate the hierarchy of the default tree.

Public Folders in Active Directory

Every public folder in a public folder store can appear as a mail recipient in the directory. You can mail-enable a public folder by selecting it in Active Directory, right-clicking it, and then clicking **Mail Enable.** After a public folder is mail-enabled, the System Attendant connects to Active Directory and creates an object for the public folder in a container such as Users. This container is specified on the **General** tab of the public folder tree properties configuration and applies to all public folders in the tree.

After a public folder is mail-enabled:

- A directory entry exists with a name: *Folder Name* + *Global Unique Identifier (GUZD).* Users with access to Active Directory can use the mail address properties of the object to send e-mail to the public folder.
- E-mail Addresses, Exchange General, and Exchange Advanced property pages are available for the public folder in Exchange System Manager (or the Exchange Folders snap-in) and the Active Directory console.
- You can configure the page to appear in the global address list for clients such as Outlook.

  Note In Exchange Server *5.5,* public folders are placed in the directory by default, but not displayed in the global address list. This changes in an Exchange 2000 mixed mode environment. New public folders are mail-enabled and configured as visible in the global address list. If you run Exchange 2000 in native mode, new public folders are not mail-enabled by default.

- The public folders list in the global address list and Exchange System Manager is managed by different Exchange Server services than the replication of public folders.

Connecting to a Public Folder

When you create a public folder, only one copy of the public folder is created by default. A public folder can exist in an organization as a single copy or as multiple copies, known as *replicas.* By using public folder replicas, you can easily provide multiple, redundant information points as well as load balancing for data access.

When a client attempts to gain access to public folder data, it must connect to a server that has a replica of the data in order to present the data to the user. For maximum efficiency, the client connects to servers in the following order:

1. The default public folder store for the client. The default public folder store is determined by the configuration of the mailbox store containing the user's mailbox. If the default public folder store is not available, the client is returned a list of servers that contain a replica.
2. If the client has an existing connection to a server in the list, it uses that server.
3. Next, it attempts to connect to each server within the same server Routing Group as the public folder server routing group for the client.
4. If it cannot connect to servers within the routing group, the local Web Storage System instructs the client to attempt a connection with other routing groups in the order of their routing group connection value.

If connections to two servers have the same routing cost, Exchange pools together the servers containing a replica and randomly selects them as if they were in the sample site.

> **Note** Cost is a means of organizing routing group connectors so that you can determine which connector is the best to use. You can assign each connector a cost to help determine load balancing. For example, suppose you have two routing group connectors: Connector A and Connector B. Connector A has a wider bandwidth and higher speed than Connector B. You might then assign Connector A the cost of 10 and Connector B the cost of 30. Because Connector A "costs" less, it is better, faster, and less troublesome to use. Because Connector B costs more than the other server, Exchange uses it less often.

The client connection method is simpler than previous versions of Exchange because Exchange 2000 does not use the location or Affinity configuration parameters. Affinity is the condition under which Exchange uses a routing group to refer the client to another server, allowing the client access to content when a specific server is not known and eliminating the need for a cost-based referral list. In previous versions of Exchange, a location identifies a subset of servers in a site that are connected by a high bandwidth network. It is needed because sites often include additional remote servers for administrative reasons. Locations help identify local servers that clients should use first when looking for a public folder replica.

In Exchange 2000, location specification is not necessary since a routing group is designed to include only servers connected with a high bandwidth network, while administrative groups resolve the administrative issues.

Previous versions of Exchange use Public Folder Affinity to enable public referrals to servers in remote sites and determine the cost/order to use for each of these sites. Exchange 2000 enables referrals between routing groups in the routing group configuration. The cost assigned to the connector between routing groups determines the cost for each remote routing group.

For more information about Public Folder Affinity, see the Exchange 2000 online documentation.

Message Flow

Message flow is the movement of a message through the Exchange 2000 messaging system. Several components handle message envelopes and their contents. The Exchange 2000 Message Routing Architecture uses a routing group master that keeps track of all changes in the messaging links and propagates link state information to all servers in the group. Exchange 2000 uses the Link State Algorithm, which determines the most efficient open route for messages, working around inoperable servers, links, or bottlenecks downstream.

The Message Transport Architecture uses a full-featured SMTP transport for its native communications while the X.400 connector remains intact for mixed-mode connections. Inbound and Outbound messages flow through SMTP and MTA for streamlining and efficiency.

Message Routing Architecture

The Exchange 2000 message routing architecture uses Active Directory, routing and administrative groups, and domains for flexibility and efficiency.

Exchange 2000 uses the following sources to accomplish its routing tasks:

- Active Directory, to determine recipient information, including information about on which mailbox a user is homed. The Active Directory Domain structure is used to provide flexibility in user and contact administration.
- Administrative groups, to delegate administrative permissions.
- Routing groups, to control routing topology and message flow. Within a routing group, all message flow is single hop point-to-point.

  Note When operating in native mode in Exchange 2000, you can completely separate administration groups and routing groups.

Routing Group Master and Propagation

In Exchange 2000, each routing group contains a routing group master, which owns the routing table for the routing group. When a routing group bridgehead malfunctions, the routing group master is immediately notified. This data is then propagated to all other servers in the routing group.

Link State Algorithm

The link state algorithm uses bridgehead availability information distributed by the routing group master to make routing decisions. The link state algorithm uses the **TCP** connection on port 691 for propagation within the routing group and X-LINK2STATE SMTP verb extension for propagation between routing groups. Link state propagation takes place in almost real time, allowing servers to make intelligent choices about message routing alternatives. This reduces message bouncing between servers as each Exchange 2000 server determines whether or not alternate or redundant links are functioning. It also overcomes message-looping problems.

The link state algorithm is similar to the Open Shortest Path First **(OSPF)** networking algorithm, which many network routers use.

Message Transport

All fast, scalable, and reliable messaging systems rely on a strong underlying transport and routing engine. This core component is fundamental to the operation of an enterprise messaging system with users who might be located worldwide. Without an intelligent transport system, messaging servers function independently of one another.

Earlier versions of Exchange incorporate a message transfer agent (MTA) built on the X.400 standard. Exchange site boundaries define the message routing topology in which it routes information with a single hop using RPC communications. You can deploy a number of connectors between sites to enable messaging, ranging from Site Connectors that used **RPC** to X.400 connectors and Dynamic Remote Access Service (RAS) connectors.

Exchange 2000 includes a modified version of the Exchange MTA but uses a full-featured SMTP transport for all native communications. The X.400 MTA is part of Exchange 2000, and it supports Exchange 2000's advanced routing engine. Exchange 2000 uses the MTA as a protocol engine to connect to earlier Exchange servers through RPC or when the X.400 messaging protocol is used.

By using SMTP as the native communication method between Exchange, you can eliminate some of the deployment issues of earlier transport methods. For example, organizations with a distributed user base normally design their Exchange site models according to the available network bandwidth rather than designing them for administrative convenience. This is because all Exchange servers within a site use **RPC** to communicate with one another, and low-bandwidth and high-latency networks are inefficient (and sometimes unworkable) for the synchronous nature of **RPC.** Because Exchange 2000 does not use **RPC** to transfer messages, you can create a more flexible routing scheme.

Component Storage Location

The Exchange 2000 message flow architecture components and their storage location in the file directory are as follows:

- IIS process (Inetinfo.exe) is located in \Winnt\System32\Inetserv
- Drviis.dll is located in \Win\System32\Inetserv

- Exsmtp.dll is located in \Exchsrvr\Bin
- Store.exe is located in \Exchsrvr\Bin
- ExIPC.dll is located in \Winnt\Systern32
- Emsmata.exe is located in \Exchsrvr\Bin
- Ntfsdrv.dll is located in \Winnt\System32\Inetsrv
- Exifs.sys is located in \Winnt\System32
- Aqueue.dll is located in \Winnt\System32\Inetserv
- Phatcat.dll is located in \Exchsrvr\Bin

Figure 26.9 shows a diagram of the components and their relation to one another.

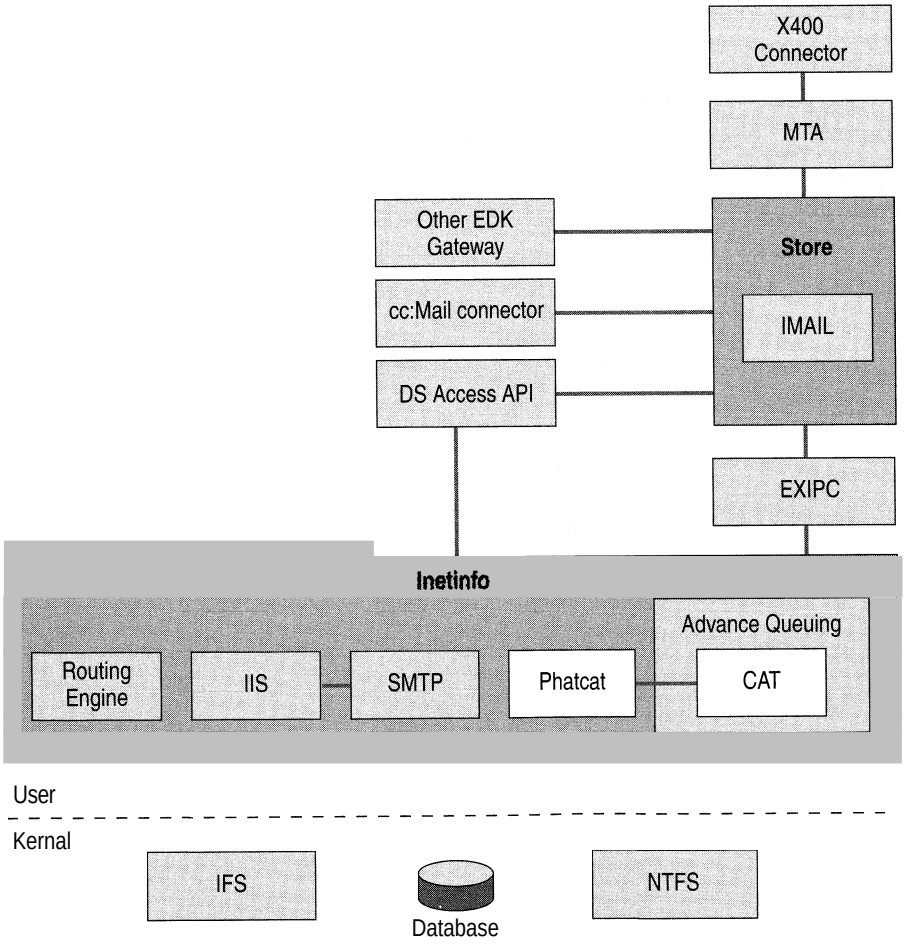


Figure 26.9 Exchange 2000 message handling architecture components

Windows 2000 Message Flow

The following steps show the basic message flow process in Windows 2000 Server:

1. The SMTP service in Windows 2000 receives the message. After negotiation, the SMTP service receives the first command dealing with the message.
2. An IMAILMSG structure is created and the SMTP service stores envelope properties into IMAILMSG. The SMTP service receives the DATA (or equivalent) command and asks the IMAILMSG for the content handle.
3. The SMTP service streams the RFC 822 message to the NTFS file system. When the message is successfully received, the SMTP service requests the IMAILMSG to commit, which guarantees that all data for this message (including both the envelope and the message data) is committed.
4. The IMAILMSG structure passes through the advanced queuing engine into the message categorizer.

   Note By default, the message categorizer is not enabled in Windows 2000. An enabled categorizer allows distribution list expansion and NTDS queries.

5. The advanced queuing engine queues messages appropriately after a query to the Domain Name System (DNS) server. In Windows 2000, there is no built-in routing engine and the fallback is to DNS routing.

   Note Various protocol events fire at points throughout the SMTP transaction and event sinks can hook the events.

Figure 26.10 shows the basic message flow process in Windows 2000 Server.

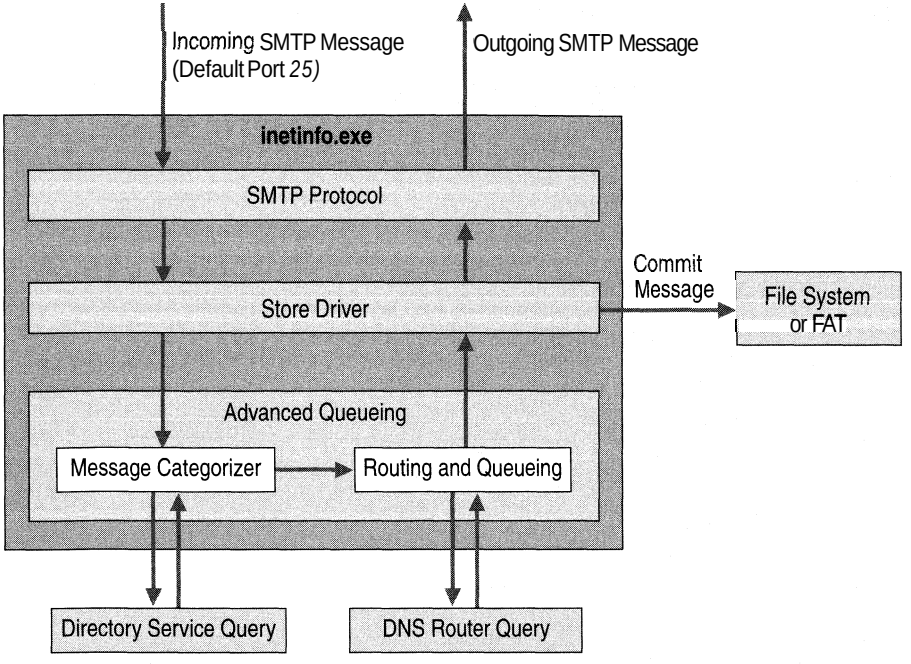


Figure 26.10 Basic message flow in Windows 2000 Server

Inbound Message Flow from SMTP to Web Storage System

The following steps show the message flow process for inbound messages from SMTP to Web Storage System:

1. An SMTP client opens a connection on the SMTP service.
2. The IIS process (Inetinfo.exe) on the SMTP service listens on port 25 and a SMTP service thread responds.
3. After negotiation, the SMTP service receives the first command regarding the message.
4. Various protocol events fire at points throughout the SMTP transaction and event sinks can hook the events. The message command is received and accepted and a store driver event is raised.
5. An IMAILMSG structure is created. The SMTP service starts storing envelope properties (for example, sender and recipients) into IMAILMSG.
6. The SMTP service receives the DATA (or equivalent) command.
7. To store the content, the SMTP service calls GetContentHandle on the IMAILMSG.
8. IMAILMSG calls the store driver to obtain the content handle and envelope handle to store the content or data.

9. The default store driver returns a file system pointer to the message body.
10. The pointer is stored as the second piece of information in the IMAILMSG.
11. The NTFS store driver writes the message body into the Queue directory.
12. Upon successful receipt of the message, the SMTP service requests the IMAILMSG to commit, which guarantees that all data for this message (including both the envelope and message data) is committed.
13. The IMAILMSG structure is handed to the advanced queuing engine, which places it in the Pre-Categorizer Queue.
14. A message categorizer thread picks up the message from the Pre-Categorizer queue for processing.

    Note This process involves several steps and consists of resolving the sender and then each recipient for the message until it resolves all recipients. If the envelope recipient list includes distribution groups (the Windows 2000 equivalent of distribution lists), it expands the recipient list to include those members. This is true if expansion for that distribution list is allowed on this server. A per-sender and per-recipient limit check is applied and recipients marked appropriately. Each recipient is marked as either Gateway (meaning that the recipient is reachable via the MTA) or Local (meaning that the recipient is actually on a mailbox store on this server). The Categorizer raises several events into which event sinks can hook.

15. On completion of the Categorization process, the message goes into a per-destination-donlain queue inside the advanced queuing engine.
16. The advanced queuing engine now passes the destination to the Routing Engine, which returns with a next-hop identifier for that destination.
17. If a message is destined for the local store, the database is stamped on the message and it is placed in the local delivery queue. The advanced queuing engine then calls the store driver to deliver the message by raising a store driver event.
18. The advanced queuing engine passes the IMAILMSG and the recipient indexes to the recipients to review.
19. The Exchange 2000 store driver picks up the message and (seeing that the recipient is marked as Local) submits it to Web Storage System using the store-side store driver interface.

Figure 26.11 shows the inbound message flow from SMTP to Web Storage System.

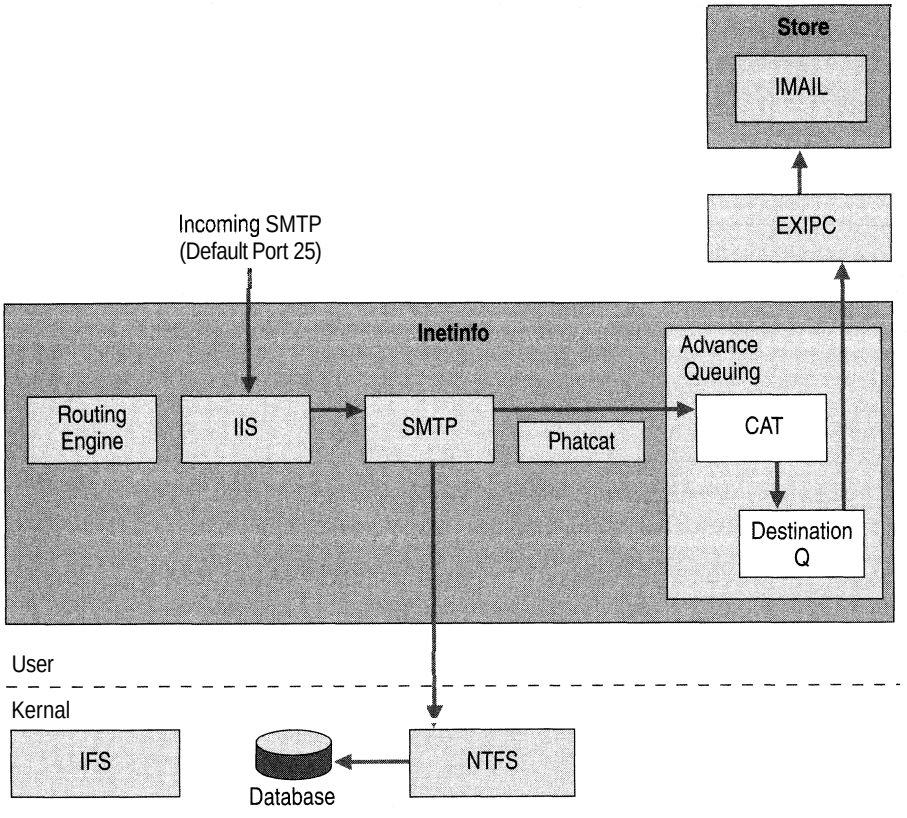


Figure 26.11 Inbound message flow from SMTP to Web Storage System

Outbound Message Flow From Web Storage System to SMTP

The following steps show the message flow process from Web Storage System to SMTP:

1. A message is submitted by a MAPI client.
2. Web Storage System makes queries to the DS/DA Access API and applies quota limits on sending.
3. Web Storage System puts the message into its SendQ folder and raises an event for the store driver.
4. The store driver picks up the message. The store driver constructs an ImailMsg and gives that to advanced queuing engine. (The message body at this time is a MAPI body).
5. A message categorizer thread picks up the message from the pre-categorizer queue for processing.

Note This process involves several steps and consists of resolving the sender and then each recipient for the message until it resolves all recipients. If the envelope recipient list includes distribution groups (the Windows 2000 equivalent of distribution lists), it expands the recipient list to include those members. This is true if expansion for that distribution list is allowed on the server. A per-sender and per-recipient limit check is applied and recipients are marked appropriately. Each recipient is marked as either Gateway (meaning that the recipient is reachable via the MTA) or Local (meaning that the recipient is actually on a mailbox store). The categorizer raises several events into which event sinks can hook.

6. The advanced queuing engine passes the destination to the routing engine, which returns with a next-hop identifier for that destination.
7. When the categorization process completes, the message is put into a per-destination-domain queue inside the advanced queuing engine.
8. If a message is destined for the local store, the MDB is stamped on the message and it is placed in the local delivery queue. The advanced queuing engine then calls the store driver to deliver the message by raising a store driver event.
9. If a message is destined for another server via SMTP, it is placed in a queue associated with its final destination identifier.
10. The advanced queuing engine makes a call into the routing engine, which returns a next-hop identifier, which gets matched up with a link to that next-hop. The link is an SMTP connection to a particular next-hop identifier (either a bridgehead group or a server).

Figure 26.12 shows the outbound message flow from Web Storage System to SMTP

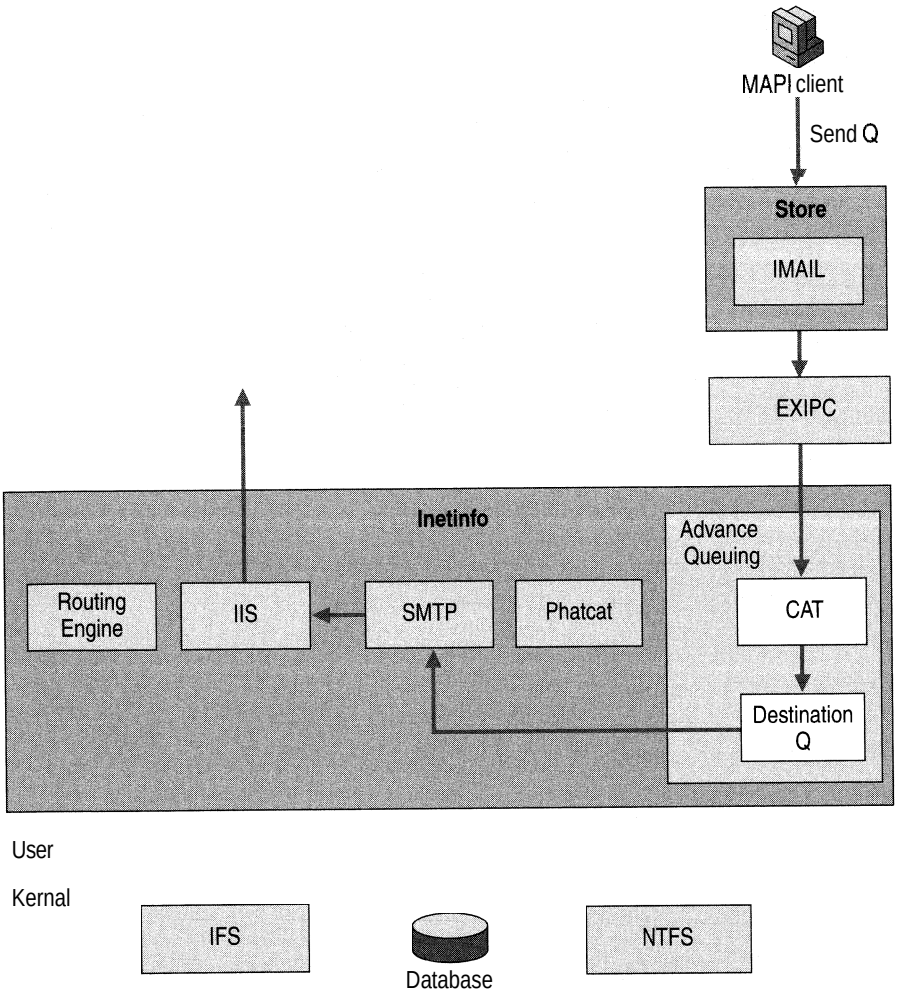


Figure 26.12 Outbound message flow from Web Storage System to SMTP

Inbound Message Flow Through the MTA

The X.400 inbound message flow is very similar to previous versions of Exchange. The difference is that once the message comes into the MTA, it always passes the message back to advanced queuing within the SMTP service for routing. The MTA simply puts the message in the SMTP MTS-Out Store folder. The remaining part of the process is exactly as for SMTP, since the store does not handle the message any differently than other messages. For more information, see "InBound Message Flow from SMTP to Web Storage System" earlier in this chapter.

Figure 26.13 shows inbound message flow through the MTA.

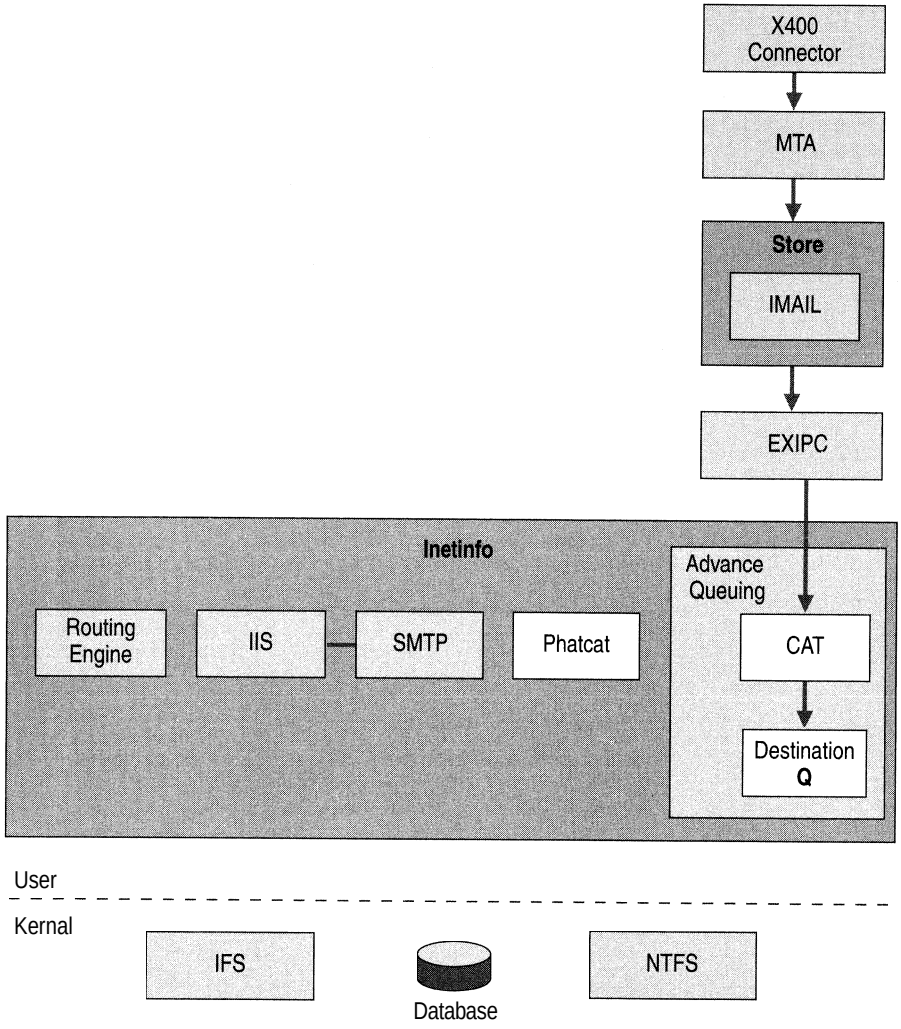


Figure 26.13 Inbound message flow through the MTA

Outlook Web Access Architecture

Outlook Web Access in Exchange 2000 does not use MAPI to communicate with the mailbox store, and Active Server Pages (ASP) are not used for client access. Clients access the Web Access Component using HTTP; however, the IIS server hosting the component uses Web Storage System to provide access to the user's messaging functions.

IIS receives Outlook Web Access client requests as a proxy for message traffic between a Web client and an Exchange server. When IIS receives a client request, it looks at the namespace and passes the client to the appropriate Outlook Web Access scripts registered for that URL. If the server contains the Exchange 2000 database, Outlook Web Access uses a high-speed channel to access the mailbox store. If the server is a front-end server, Outlook Web Access proxies the request to a back-end server using HTTP. The following sections describe Outlook Web Access features, components, and processes.

Accessing a Server

When using a client program such as Outlook, the user interacts directly with the Exchange server. However, with Outlook Web Access, the user interacts with the IIS Web service from the browser. The browser communicates with the server by using HTTP and Web Distributed Authoring and Versioning (WebDAV).

When IIS receives a client request for an item in a virtual directory mapped to Web Storage System, it gives the request to an Exchange Internet Services Application Programming Interface (ISAPI) application, which communicates with Web Storage System. Web Storage System returns the requested data and the ISAPI application renders it into the appropriate HTML for the client's browser.

In addition to HTML, Outlook Web Access sends additional data to Microsoft Internet Explorer 5 clients by using Extensible Markup Language (XML). This allows the client to do more processing and send fewer requests to the server.

In a more scalable or distributed environment, one or more front-end servers can process a client's requests and route them back to the back-end server that contains the client user's mailbox.

Outlook Web Access uses HTTP and WebDAV to communicate between client browsers and the Outlook Web Access server. In large sites, you can place user mailboxes on multiple back-end servers that are referenced by one or more front-end servers. This multiple server architecture provides additional scalability and a single namespace for back-end servers.

Clients direct specific requests to Outlook Web Access by using named URLs. Often the URL, such as http://owa.microsoft.com/exchange/username, directs the client to the user's mailbox. Named URLs can be used for more than addressing a mailbox, however. You can address most functions and components of the client by defining a specific URL.

You can open specific folders by entering the name of the folder after the mailbox name. For example, to open a calendar, type the path to the user's mailbox, then type **/calendar** (http://owa.microsoft.comlexchangeljuserlcalendar). Likewise, you can access the Contacts folder directly by entering the path to the client's mailbox followed by **/tasks.**

Named URLs can be used for more than accessing folders. You can open any item and perform many functions using explicit URL addressing. Table 26.3 shows the many option and command verbs you can use, giving you a wide range of actions.

Table 26.3 The Outlook Web Access option and command verbs

Option	Description
Page=x	Displays the navigation bar
View=*x*	Uses the Outlook view named x
Sort=*x*	Sorts by column x
Date=yyyymmdd	Displays the date in Calendar
Cmd=action	Performs the stated function

Action	Description
Navbar	Displays the navigation bar
Contents	Displays the contents of a folder
New	Creates a new default item in a folder
New&Type=x	Specifies the type of item to create
Options	Sets options
Open	Opens a message or appointment for reading
Edit	Opens a message or appointment for editing
Reply, ReplyAll, Forward	Performs message operations
Accept, Decline, Tentative	Performs appointment operations

URL usage follows this syntax:

```
http://server_name/virtual_root/folder/?option= Modifier
http://server_name/virtual_root/folder/item_name?Option=&Modifier
```

For example, the URL entered for *juser* to create a new message might be:

```
http://owa.microsoft.com/exchange/juser/?Cmd=new
```

Server Components

The Outlook Web Access server is a proxy for all message traffic in an environment where you use Web browsers to access data on an Exchange 2000-based computer. Client requests are received by the IIS 5.0 Web service and given to the Outlook Web Access ISAPI application for processing. If the server contains an Exchange 2000 database, Outlook Web Access uses a high-speed channel to access the store. If the server is a front-end server, Outlook Web Access proxies the request to a back-end server using HTTP.

> **Note** Unlike Exchange Server 5.5, IIS is a required component of Exchange 2000 and is present on every computer running Exchange 2000.

If the client uses Internet Explorer 5, the Dynamic HTML (DHTML) features of Internet Explorer perform more of the rendering on the client, which improves server performance. With DHTML, commonly used HTML and script is encapsulated and downloaded only once to the client. For all other clients, such as Internet Explorer 4.*x* and other Web browsers, most of the rendering occurs on the server and only a small amount of JavaScript goes to the client.

Log On and Mailbox Display

With Exchange 2000, you can use user logon credentials to automatically open a mailbox or you can use a URL to specify the mailbox to open. For example: http://Outlook Web Access *Servername*/exchange/*mailbox name.*

The following steps describe the flow of information when a user logs on to a mailbox and views the Inbox:

1. The user requests the Exchange 2000 mailbox by specifying the following URL in the browser: http://Outlook Web Access Servernamelexchangelmailhox.
2. The user is authenticated by the IIS Web server, which determines a user's Windows 2000 account.
3. The mailbox location for the user is queried from Active Directory. If the mailbox is on another server, the browser is redirected to that server and user authentication occurs again.
4. Outlook Web Access returns a default page for the mailbox that contains a navigation bar in the left-hand frame and a view of the mailbox contents (\exchange\mailbox name) in the right-hand frame.

Note In Exchange Server 5.5, Outlook Web Access uses a separate logon page for users. The logon page displays initially and prompts the user for the mailbox name. The user then provides authentication credentials to access the mailbox.

Logout

Exchange 2000 Outlook Web Access does not use cookies. To end your session, you must close all open instances of the browser. With Outlook Web Access in Exchange Server 5.5, the user can click a logout button that ends the user's session. It does this by invalidating a cookie that monitors the session.

Opening an Item

The following steps describe the Outlook Web Access open and display process for items such as folders and e-mail messages. To open and display an e-mail message, for example:

1. The browser sends a request for an e-mail message.

   The browser issues a GET request for a URL such as http://server/vroot/user/folder/message.eml. This URL does not have any query strings attached, which would be processed first, so the server returns a rendering of this resource based on its Message-Class and the default action configured for this class.

2. Exchange ISAPI processes the request.

   When IIS receives the request, it is passed to the Exchange ISAPI component DavEx.dll. This component parses the request for the following information and then sends the request to the Exchange store. Table 26.4 shows the passed item and its purpose.

3. Web Storage System determines the item type.

   The server verifies that the user has access to the item, determines the type of object (folder, message, task, and so on) and returns the item type and its state (read, unread, and so on) to the ISAPI application.

4. Exchange ISAPI selects the form.

   The ISAPI application takes these object attributes and looks for a form definition in the forms registry that matches the object's type. If a matching form definition is not found, a default form stored in Wmtemplates.dll is used. If the browser language is not English, language specific strings are loaded from other template libraries in the \Exchsrvr\Res\Directory.

5. Web Storage System retrieves data for the form.

   Once a form definition is found the ISAPI application parses the form, calling Web Storage System, to retrieve the data it references.

6. Exchange ISAPI renders the form.

   When the data is returned from Web Storage System. the form is rendered into the appropriate HTML and XML and goes to the client.

The HTML information is not browser or platform dependent. Outlook Web Access renders the HTML based on a variety of factors including the browser version. Non–Internet Explorer browsers receives HTML code that conforms to the HTML 3.2 standard. Internet Explorer 5 and later browsers receive DHTML, which means different elements respond to user clicks and do not require communication with the server.

Table 26.4 DavEx.dll passed items and usage

Passed Item	Used To
HTTP User-Agent Field header	Determine the browser type, version, operating system, and how to render content.
HTTP Accept-Language header	Determine the language for the rendered content.
HTTP Translate header	Determine if the content should display in a browser or if it should return without rendering to a WebDAV application such as Word 2000.
Query String	Determine a specific action to perform.

Front-End and Back-End Servers

A front-end server is an Exchange 2000 server that does not host a database, but instead forwards client requests to the back-end server for processing. The front-end server uses Lightweight Directory Access Protocol (LDAP) to query Active Directory to determine which back-end server hosts the user's mailbox. A back-end server is an Exchange 2000 server that maintains at least one database. This division of function between two servers provides several benefits, particularly in a Web environment, as described in the following sections.

- **Single Namespace** As multiple back-end servers are configured to handle additional mailboxes, it is desirable to refer to all of the servers with a single name. You can refer to a front-end server with a single name and it can proxy user requests to the correct back-end server containing that user's mailbox.

  If multiple front-end servers are configured to manage a high load of requests, a single namespace for these servers is maintained simply by configuring the DNS with one name mapped to the IP address on the server. It does not matter which front-end server the client connects to.

- **Offload SSL** Encrypting and decrypting message traffic uses a lot of CPU cycles. A front-end sewer can perform encryption work, giving the back-end server more cycles to manage Web Storage System.

- **Public Folder Referrals for IMAP4 Clients** Many IMAP4 clients do not support referrals. With this architecture, the front-end server can retrieve public folders that exist on a server other than the user's e-mail server.
- **Server Location** You can place the Exchange 2000 back-end servers that contain the databases behind a firewall for increased protection. You can configure the firewall to only allow traffic from the front-end server.

  Note Front-endhack-end architecture is available for HTTP/WebDAV, POP3, and IMAP4 clients.

Outlook Web Access Functions

Outlook Web Access allows users to perform many different functions using an Internet browser as the client program.

- **Light Messaging** Outlook Web Access provides an alternative to the full Outlook 2000 client. There are instances when a full client is either not required or not practical. Outlook 2000 is better suited to low bandwidth network usage. Outlook Web Access clients have less network traffic. For this reason, using Outlook Web Access makes sense in configurations that cannot spare bandwidth to support the full Outlook client.
- **Roving User Support** You can support roving users by using system policies and server-based profiles that carry little network traffic in terms of administration and performance. In instances where it is impractical to use Microsoft IntelliMirror features to support roving users, you should use Outlook Web Access because there is no client to install other than the browser, and no MAPI profile to consider.
- **Kiosks** You can use kiosks to place computers in locations such as factory floors, common areas, and conference rooms. This provides users access to e-mail, calendaring, and other basic messaging functions. This is an appropriate application to provide general access to posted public folders.
- **Migration** Outlook Web Access provides a good interim client access solution, which mitigates the operational impact normally associated with extended migration periods. During migration, you can move small groups at a time, and deploy custom client configurations to achieve a period of coexistence. Outlook Web Access can solve potential problems that might arise from client software coexistence.

WebDAV

With increased focus on Internet standards and network interoperability, Web Distributed Authoring and Versioning (WebDAV) is an important communication protocol for the Internet as an extension to HTTP. The WebDAV specification was written by the Internet Engineering Task Force (IETF), with significant contributions by Microsoft. WebDAV has found application in many areas and provides a Web interface to Exchange 2000 Web Storage System.

WebDAV is used in many different applications. It enables collaborative publishing to IIS servers through the Web. It is the protocol that supports Webfolders in Microsoft Office 2000. WebDAV provides a Web interface to Web Storage System in Exchange 2000 by allowing access to the hierarchical database in Exchange 2000.

Because of its inherent integration with XML, WebDAV is an excellent medium for communicating XML data over the Web. However, before the strength of combining these two technologies can be fully understood, it is important to understand what WebDAV is and how it could be useful in your Exchange 2000 client/server architecture.

WebDAV Process

Figure 26.14 shows the process by which clients access items in their Inbox using WebDAV

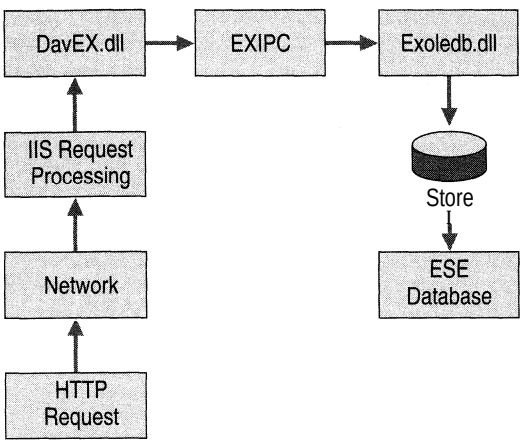


Figure 26.14 WebDAV process for accessing items in Inbox.

The following process shows how clients access items in their Inbox using WebDAV.

1. The client issues an HTTP GET request for the client's inbox.
2. IIS receives the request on port 80 (unless you change this configuration) and sends the request to the Exchange 2000 ISAPI application, DavEX.dll, for processing using ExIPC.
3. The request is forwarded using ExIPC to Web Storage System driver, Exoledb.dll.
4. Exoledb.dll renders the request in a format that Web Storage System can process. Web Storage System accesses the Extensible Storage Engine database to retrieve the client's Inbox properties.
5. Once the clients Inbox properties are retrieved, Exchange 2000 routes the information back to the client using the same components that it used to process the client request.

Exchange 2000 and Previous Versions of Exchange

Exchange 2000 can coexist with previous versions of Exchange when you run Exchange in mixed mode. Each version of Exchange can detect other versions and share data with Active Directory. Exchange 2000 uses Active Directory and previous versions of Exchange use the Exchange directory. You can set up coexistence in organizations that have multiple divisions that are managed under different management structures with different implementation schedules. Implementing Exchange 2000 requires time and planning. Coexistence is a viable solution if you want to use Exchange 2000 features immediately, before you finish migrating.

In a mixed mode Exchange organization, the procedures for managing earlier versions of Exchange apply to Exchange 2000 as well; for example, all servers on the site must use a common Site Services Account. Unlike earlier versions of Exchange, you do not need to upgrade the bridgehead servers of a site before you deploy Exchange 2000.

For more information about coexistence and upgrading, see the *Microsoft Exchange 2000 Server Planning and Installation Guide*.

Architectural Improvements in Exchange 2000

Microsoft Exchange 2000 Server is architecturally different from earlier versions of Exchange. The components of Exchange 2000 provide a more flexible, integrated. and scalable architecture than earlier versions. If you have an existing Exchange system, whether it is based on Exchange 4.0, Exchange 5.0, or Exchange 5.5, it is important to understand the coexistence and upgrade options before installing Exchange 2000.

Exchange Server 5.5 is minimally reliant on the underlying Microsoft Windows domain structure and, because of this, deployment and support personnel might have little experience with Windows operating systems. Because Exchange 2000 tightly integrates with Active Directory, you need to understand the prerequisites and dependencies of Active Directory before beginning to design your Exchange 2000 deployment.

Figure 26.15 shows a graphical representation of Exchange 5.5 Server architecture with its underlying components, including the dedicated Exchange directory and the Jet database.

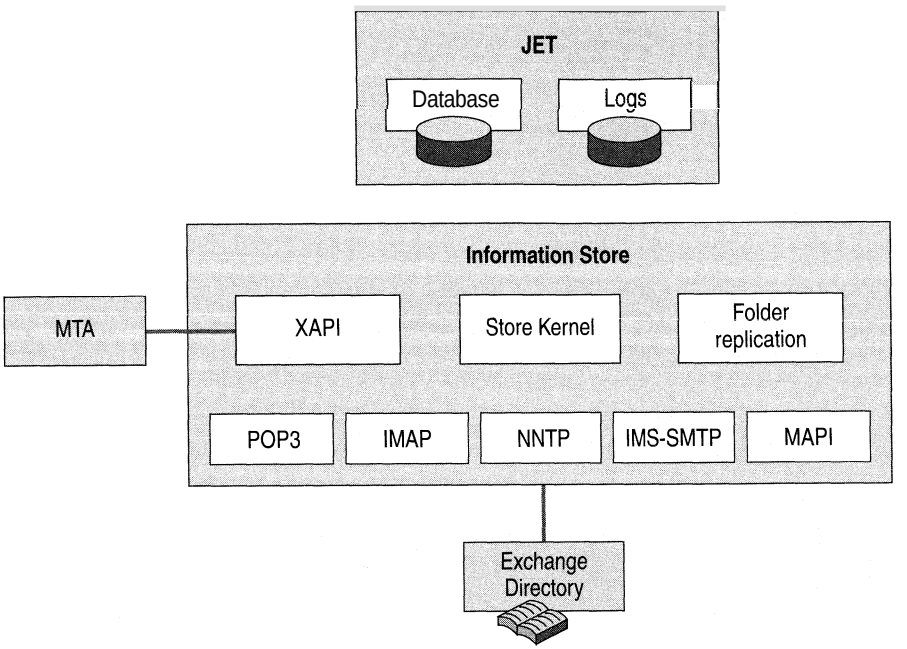


Figure 26.15 Exchange Server 5.5 architecture

Web Storage System Changes

Web Storage System in Exchange 5.5 typically maintains messages in Message Database Encapsulated Format (MDBEF). When a non-MAPI client requests a message, Web Storage System uses the IMAIL process to convert the contents from MDBEF to the appropriate MIME or non-MIME encoded format based on the client's request.

This conversion process consumes time and system resources. To reduce the network traffic, Web Storage System in Exchange 2000 includes two types of storage files: Native Content slore (.stm file) and Rich-Text store (.edb file).

Clustering Improvements

In Exchange 5.5, you can use up to two nodes in a cluster in an active/passive configuration. That is, only one instance of the Store.exe file runs (on the active node) while the other node is used strictly for fail over. If the active node fails, all storage groups are failed over to the other node. The problem with this configuration is that it consumes time and resources to an extent that are often unacceptable.

Exchange 2000 Integration with Active Directory

Unlike previous versions of Exchange, Exchange 2000 doesn't have an integrated directory. Instead, Exchange 2000 integrates with Active Directory. Unlike the Microsoft Windows NT Security Accounts Manager (SAM), which is not designed to hold information about directory objects, such as telephone numbers, addresses, and certificates, Active Directory can hold the directory information required by Exchange 2000. Integration with Active Directory provides increased system performance and manageability, and makes directory management easier. Some of the features of Active Directory include:

- **Centralized object management** Unified administration of Exchange 2000 and Windows NT directory objects allows you to manage all user data in one place, with one set of tools.
- **Simplified security management** Native Windows 2000 Access Control Lists (ACLs) are used in Exchange 2000 Web Storage System, so a single set of security groups administered once applies to data stored in Exchange 2000 and Windows 2000 file shares.
- **More efficient creation of distribution lists** You can use security groups in Windows 2000 as distribution lists, which removes the need to create a parallel set of distribution lists for each department or group.
- **Easier access to directory information** You can increase flexibility of the overall enterprise directory solution through schema extensibility for management of distributed information, use of LDAP as a native access protocol for directory information, and easy hierarchy reconfiguration.

All Exchange 2000 directory information (including mailboxes, information about servers and sites, and custom recipients) is stored in Active Directory. Distribution lists are based on Active Directory groups, simplifying list administration.

Active Directory Connector

Recognizing that customers might migrate to Exchange 2000 over time, Exchange 2000 provides Active Directory Connector (ADC), which you can use to replicate directory information between Exchange 2000 and existing Exchange 5.5 sites. ADC consists of a replication and mapping engine in addition to engines for both Exchange *5.5* and Active Directory. ADC provides a mechanism for replicating an Exchange *5.5* directory with Active Directory.

ADC uses administrator-defined replication agreements, which define the direction of replication, schedule for replication, and placement of replicated objects.

ADC works in the following ways:

- Uses the LDAP API to perform fast replication between the Exchange 5.5 and Windows 2000 Active Directory directories.
- Hosts all active replication components on Active Directory.
- Replicates changes only between the two directories.
- Maintains object fidelity through replications (for example, Active Directory Group object to Exchange Distribution List object).
- Hosts multiple connections on a single Active Directory server and manages these through connection agreements.

Exchange 2000 does not have a directory service of its own; instead, it relies on Active Directory, which is provided by Windows 2000, for browsing, security, and name resolution. For organizations that use Exchange 5.5, coexistence between the Exchange 5.5 directory and Active Directory is very important before upgrading to Exchange 2000. Figure 26.16 shows the Active Directory Connector architecture.

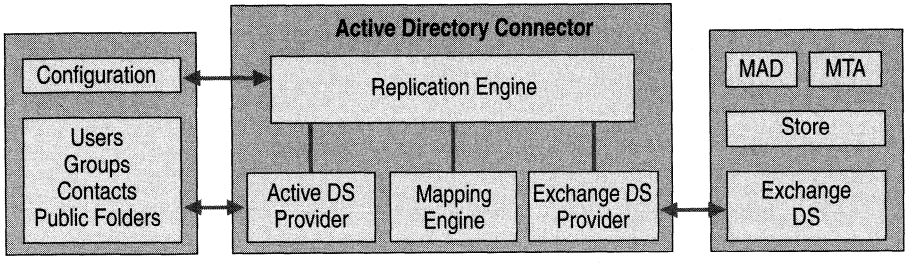


Figure 26.16 Active Directory Connector architecture

Active Directory Connector Process

When a mixed mode connection is required, the Active Directory Connector process is initiated. The following sections describe each step that ADC takes when replicating files from an Exchange 5.5 environment to an Exchange 2000 environment.

Initial Replication

When the connection agreement for the ADC is initially executed, the stored Update Sequence Number (USN) on the connection agreement is set to 0; thus, ADC functions as if the agreement is run for the first time. Each connection agreement has its own signature, which computes when the connection agreement configures. Active Directory objects replicated into the Exchange directory have the same DSA-Signature attribute as the Exchange 5.5 server bridgehead, although the Replication-Signature attribute of the object matches the computed value of the connection agreement's signature. When an Active Directory object replicates into the Exchange directory, the Object-Version attribute, which is standard to Exchange, is either set to 1 (if it is a new object), or incremented by 1 (if it is modified).

Before the LDAP write is made, the current value of the Object-Version (if it exists) is read, incremented by 1, and then written into the Replicated-Object-Version attribute, also on the Exchange directory object. Thus, if ADC modifies an object, both the Object-Version and Replicated-Object-Version attributes are the same.

Detecting Changes in the Exchange Directory

When an LDAP search to the Exchange directory occurs, the request is for objects from the stored USN (on the connection agreement) to the highest USN in the Exchange directory. This also includes entries of deleted objects. A filter is set so that objects that have the signature of the connection agreement are not requested unless the Object-Version attribute is greater than the Replicated-Object-Version attribute; this prevents ADC from back-replicating unchanged objects previously sent to the Exchange directory.

Detecting Changes in Active Directory

Active Directory uses attribute-based replication, unlike the Exchange directory, which uses object-based replication. To detect changes in Active Directory that you need to replicate to the Exchange environment, the connection agreement uses a combination of Active Directory USNs and the sum of Attribute Versions of each Active Directory object in the source container.

Object Class Mapping and Attributes

An Exchange 5.5 mailbox object mapped to a Windows NT 4.0 account appears in the Exchange directory as a Mailbox object but maps to Active Directory as a disabled Windows user. You can change this default behavior of the connection agreement to map it as either an enabled user or as a mail-enabled contact in the configuration interface for the connection agreement.

> **Note** You can prevent ADC from creating an object altogether if the Exchange directory object cannot map to Windows 2000 Active Directory.

An Exchange 5.5 mailbox object mapped to an Active Directory account appears in the Exchange directory as a mailbox object, and maps to Active Directory as a mailbox-enabled user (the *msExchHomeServerName* attribute is set). The Object-GUID attribute of the Exchange mailbox object is set to the globally unique identifier (GUID) of the Active Directory object, and the *legacyExchangeDN* attribute of the Active Directory object is set to the distinguished name of the correlating object in the Exchange directory. All directory attributes from the two objects, such as telephone number, postal address, and so on, are merged and populated to both directory objects.

A distribution list object in the Exchange directory appears in Active Directory as a mail-enabled Group object (type: Distribution, scope: Universal). Because the distribution list object can appear in Active Directory before the membership objects exist, any orphaned members are binary encoded and written to the *unmergedAtts* attribute in the corresponding entry of the group object. This ensures that membership changes to the Active Directory group object successfully replicate back to the Exchange directory. The unmerged attributes are removed and resolved only after a full replication of the object is initiated.

A custom recipient (of any address type) in the Exchange directory appears as a mail-enabled contact in Active Directory.

Table 26.5 lists the objects that replicate from Active Directory to the Exchange directory.

Table 26.5 Objects that replicate from Active Directory to the Exchange directory

Active Directory Object	Exchange Object Class Mapping
Mailbox-enabled user	Mailbox
Mail-enabled user	Custom recipient in the target container
Non-mail-enabled user	Not replicated
Mail-enabled contact	Custom recipient in the target container
Non-mail-enabled contact	Not replicated
Mail-enabled group (type: Distribution)	Distribution list in the target container
Mail-enabled group (type: Security)	Distribution list in the target container
Non-mail-enabled group (type: Distribution)	Not replicated
Non-mail-enabled group (type: Security)	Not replicated

Duplicate Object Detection

If at any point ADC attempts to create an object that already exists in the target directory, a numeric value, prefixed with a hyphen, is appended to the common name of the object. The prefix ranges from -1 to -9999.

Schema Discovery

When replicating data between two different systems, the format and restrictions for the data might be different for each system; for example, Active Directory supports UTF8 in the object name, but the corresponding entry in the Exchange directory—in this case, the distinguished name—supports only plain text characters. Schema discovery allows ADC to work out the restrictions imposed by the target directory and perform the necessary conversion. Schema discovery is designed to accommodate the following discrepancies:

- Data format and presentation (for example, UTF8 and teletext)
- Field length restrictions
- Mandatory and optional attribute mapping
- Single-value to multi-value field mappings (and vice versa)

Exchange 2000 Dependencies on Windows 2000 and Active Directory

Exchange 2000 relies on Windows 2000 in the following ways:

- Active Directory becomes the Address Book for Exchange and Outlook clients.
- Active Directory is the configuration repository for your Exchange organization.
- Exchange 2000 uses and extends Windows 2000 SMTP, NNTP, and HTTP protocol stacks.
- Exchange 2000, like Windows 2000, uses DNS for name resolution.

Earlier versions of Exchange integrate with the Windows NT security model, but directory information is contained in a separate database (Dir.edb). Because of its architecture, for Exchange 2000 to fully coexist with earlier versions of Exchange, replication between Active Directory and the Exchange 5.5 directory must take place. ADC coordinates replication.

CHAPTER 27

Application Development

A collaboration application is a program that facilitates groups working together by collecting, organizing, distributing, and tracking information across an organization. When an application designed for personal use is placed in a public folder, it becomes available to others. An effective collaboration environment streamlines workflow so colleagues can interact efficiently, find and share information, collaborate on documents, and publish information to the company intranet or the Internet. The growing integration of Web-based technologies to build collaboration applications requires that you be more involved to ensure the reliability and integrity of mission critical operations. By understanding development terminology and basic programming procedures, you can better help deploy applications while maintaining server policies and security.

This chapter discusses the development capabilities of Microsoft Exchange 2000 Server and Microsoft Web Storage System. Messages, files, Web content and semi-structured data can all be stored in Web Storage System. Web Storage System provides easy access to its contents and versatile storage for developers' applications using Web protocols. The many protocols used by Exchange make it an ideal vehicle for developing applications. It is this versatility that makes Microsoft Exchange 2000 a rich development platform.

In This Chapter

Overview of Web-Based Collaboration

Exchange 2000 Web-Based Collaboration Applications

Event Support in Exchange 2000

Workflow Components for Exchange 2000

Application Development Tools

Overview of Web-Based Collaboration

Microsoft Exchange 2000 Server provides the infrastructure for building collaboration applications. Exchange 2000 uses a variety of protocols, application programming interface (API) components, and formats to permit communication between applications and Web Storage System.

Before reviewing how Exchange 2000 Server works, it is important to know what you can accomplish with Exchange 2000. For example, Exchange 2000 technology allows you to construct applications specific to your organization that can display functional in-boxes, display corporate information on Active Server Pages (ASPs), while allowing users to track expense reports and other business data.

Another example is an Exchange 2000 application called Microsoft Transit, which demonstrates how you can use various features of Web Storage System to create a career management tool. You can place Transit in the Human Resources space of a company and it can use career maps created by users to find personal and collaborative job spaces. In personal job spaces users can create goals, respond to employee review requirements, collaborate with peers and direct reports, and explore career options within the company. In public job spaces, users can interact with others who have similar professional interests, create networks, and find professional development information.

Exchange 2000 uses the following protocols, components, and formats:

Network Protocols

- Hypertext Transfer Protocol (HTTP)
- World Wide Web Distributed Authoring and Versioning (WebDAV)
- Simple Mail Transfer Protocol (SMTP)
- Network News Transfer Protocol (NNTP)
- Internet Message Access Protocol version 4 (IMAP4)
- Post Office Protocol version 3 (POP3)
- Lightweight Directory Access Protocol (LDAP)

API **Components**

- ActiveX Data Objects (ADO)
- Active Directory Service Interfaces (ADSI)
- Collaboration Data Objects (CDO)
- OLEDB
- MAPI

Formats

- Request for Comments 822 (RFC 822)
- Multipurpose Internet Mail Extensions (MIME)
- HTML
- Extensible Markup Language (XML)

Only processes running on the server can use the OLE DB provider. To access the Web Storage System contents remotely, you must use the protocols listed in the following section.

Network Protocols

Exchange 2000 network protocols are not located in Web Storage System. They run as part of the Microsoft Internet Information Services (IIS) process. The advantage of removing the protocols from Web Storage System is that it allows you to dedicate servers to specific tasks, such as managing the database or handling client requests. The removal of the protocols from Web Storage System increases the flexibility, reliability, and scalability of Exchange 2000.

HTTP and WebDAV Extension

Hypertext Transfer Protocol (HTTP) is a way of connecting to the Internet. An HTTP client such as Microsoft Internet Explorer is a Web browser that uses HTTP to access information on the IIS server. Microsoft Outlook Web Access enables Exchange 2000 to integrate with an IIS server and connect to the Internet. This allows users to access e-mail and public folders in Web Storage System using a standard Web browser.

The three main HTTP functions used by Web Storage System are GET, PUT and POST:

- **GET** Retrieves a document
- **PUT** Puts an item in a folder
- **POST** Submits an item to a folder

The World Wide Web Distributed Authoring and Versioning (WebDAV) protocol provides access to an extensible set of associated properties. WebDAV also defines protocol commands used to search, move, copy, delete, lock and unlock resources, and make new collections of resources (folders). The commands **PropFind** and **PropPatch** are used for property (or attribute) manipulation. These features make WebDAV an ideal protocol for creating interoperable, collaborative applications. Major features of the protocol include:

- **Locking (concurrency control)** Long-duration exclusive and shared write locks prevent the possibility of two or more collaborators writing to the same resource without first merging changes. The duration of WebDAV locks is independent of any individual network connection. This serves two purposes: it provides robust Internet-scale collaboration, because network connections might be disconnected arbitrarily, and improves scalability, because each open connection consumes server resources.
- **Properties** XML properties provide storage for metadata, such as a list of authors or Web resources. These properties can be set, deleted, and retrieved using the WebDAV protocol.
- **Namespace manipulation** Because resources might need to be copied or moved as a Web site evolves, WebDAV supports copy and move operations. Collections, similar to file system directories, can be created and listed.

The main WebDAV functions used with Web Storage System are:

- **MKCOL** Makes a collection that is in essence the same as a Web Storage System folder. MKCOL allows you to make folders and set their properties.
- **LOCK** Locks documents to prevent writes.
- **UNLOCK** Allows writes and creates updates.
- **PROPFIND** Searches for properties individually or for a whole set of folders.
- **SEARCH** Provides filter, subquery and sort capabilities to a search.
- **PROPPATCH** Sets arbitrary properties on an item.
- **SUBSCRIBE** Adds a user to a list.
- **UNSUBSCRIBE** Removes a user from a list.
- **POLL** Checks to see if the notification is fired.

SMTP

Simple Mail Transfer Protocol (SMTP) is the Internet standard for transporting and delivering electronic messages. Based on specifications in RFC 821 and RFC 822, Microsoft SMTP Service is included in Microsoft Windows 2000 and Exchange 2000, and uses the primary SMTP address in the default recipient policy as your organization's default domain.

Exchange 2000 Server expands the SMTP Service, enhancing the basic delivery functions of the protocol without affecting its compatibility with other messaging systems. Exchange gives you greater control over message routing and delivery. and provides secure access and channels for managing the service. Although you can make some configurations for sending e-mail to remote domains on the virtual server, you should do most of the administrative work at the SMTP connector.

If you use the SMTP Service on a computer before installing Exchange 2000, any configurations you make to the operating system for SMTP are lost. Also, messages in the Pickup or Queue directories are not delivered. In Windows 2000, SMTP uses the subdirectories in the <root>\Inetpub\mailroot directory. Although these folders aren't deleted, SMTP uses the <root>\Exchsrvr\mailroot directory.

Multiple virtual servers can be helpful in certain situations. If you have different groups of users with varying security requirements or message-size needs, you might want to create additional virtual servers. Additional virtual servers are also helpful for managing different types of e-mail. For example, in a mail gateway, one dedicated virtual server can handle Internet e-mail, while another handles internal e-mail.

NNTP

Network News Transfer Protocol (NNTP) is an application protocol that is used over TCP/IP networks. NNTP defines a set of client and server commands used to access newsgroups.

Exchange 2000 Server uses NNTP to enable Microsoft Outlook 2000 users to participate in online discussions over the Internet. Exchange also enables users running client applications that support NNTP to access newsgroup public folders on computers running Exchange.

Users can read and post items, such as messages and documents, to NNTP newsgroups that are represented in Exchange as public folders; for example, scientists can exchange research information by posting messages to a newsgroup public folder for their area of interest. Other scientists worldwide can read and respond to items in the newsgroup. Items in newsgroups can be replicated to USENET host computers through news feeds.

IMAP4

Internet Message Access Protocol version 4 (IMAP4) is an Internet messaging protocol that enables a client to access e-mail on a server instead of downloading it to the user's computer. IMAP4 is designed for an environment in which a user logs on to the server from several different workstations. In such an environment, downloading a user's mail to a specific computer is usually impractical because the user does not always use the same computer. For example, IMAP4 is widely used in universities where students connect to the mail server from different labs throughout the campus. Once connected, users can access their mailboxes as though their mail is stored locally. IMAP4 does not provide mail transport. Simple Mail Transfer Protocol (SMTP) provides this feature.

IMAP4 allows a client to access private messages and public folders on a server. Users with an IMAP4 client can access mail in their Exchange mailboxes without downloading the entire mailbox to a specific computer.

A single client can access multiple mailboxes to retrieve specific messages or portions of a message, such as an attachment. IMAP4 clients can also search a mailbox and store flags to identify messages that have been read.

You can configure your IMAP4 server to grant or deny access to specific computers, groups of computers, or domains. You can grant or deny access to a single computer based on an IP address, or override IMAP4 access on a per-user basis. A group of computers can either receive or be denied access based on their subnet address and subnet mask. You can also control access to an entire domain by specifying the domain name.

POP3

Post Office Protocol version 3 (POP3) is an Internet protocol that allows a POP3 client to download e-mail from a server. This protocol works well for computers that cannot maintain a continuous connection to a server.

Both IMAP4 and POP3 are Internet messaging protocols that allow clients to access their mail. The difference between these protocols is the location where the client manipulates the messages. IMAP4 allows a client to access and manage mail on a server. POP3 allows a client to download e-mail from an Inbox on a server to the client computer, where messages are managed.

Unlike IMAP4, POP3 does not allow users to manipulate messages on the server. Mail simply downloads messages to the client where they are managed. POP3 provides access to a user's lnbox only; it does not support access to public folders.

All management of POP3 user mailboxes is performed in the Microsoft Active Directory directory service. In previous versions of Exchange, Exchange used its own database to hold directory information, and managed recipients through the Exchange administrator. In Exchange 2000, Active Directory can hold all messaging directory data. Because a separate directory for Exchange is not necessary, all mailbox administration is coordinated in Directory Service Manager. After installing System Manager on a computer running Windows 2000, a set of extensions is added to the standard Active Directory console. This allows an Exchange mailbox to be created when a new user account is defined in Active Directory Users and Computers.

Only users with Windows 2000 accounts can have a mailbox, and send and receive mail. If an account is mail-enabled but not mailbox-enabled, users can only send mail. Mailbox-enabled users can also receive mail and configure additional settings.

A unified namespace provides easier administration of multiple POP3 servers. For example, if you have three separate computers running a POP3 virtual server, you normally divide the user load by configuring certain users to connect to POP3Server1, other users to connect to POP3Server2, and still other users to connect to POP3Server3. If all POP3 servers are part of a front-endhack-end configuration, there is one name that provides user access to all POP3 servers in your configuration. You can configure clients to connect using the same POP3Server. You can use software load balancing or hardware load balancing to randomly distribute the load to any of the three POP3 servers. When you want to move a user's mailbox from one server to another, the client does not need to reconfigure the name of the server it logs on to. As your user population grows, you can add more computers to the front-end bank of servers without reconfiguring the clients.

When connections are made using Secure Sockets Layer (SSL), information is encrypted and decrypted. The encryption and decryption process is processor intensive and can affect performance. If your POP3 virtual servers are deployed in a front-endhack-end configuration, the front-end servers can process encryption with the client. When the front-end server and back-end server communicate, they do so without the network bandwidth load of SSL encryption. This reduces the load on the back-end server.

LDAP

The Lightweight Directory Access Protocol (LDAP) protocol is a distributed, hierarchical directory service protocol that you use to gain access to repositories of users and other network objects. Because LDAP is not typically tightly integrated with the host operating system, information can be kept in both LDAP and in a name service such as Network Information Service. When you use LDAP to gain access to the groups, you reduce redundancy and maximize LDAP's scalability.

LDAP gives you the option of using paged or non-paged results when performing a search. With non-paged results, the maximum number of objects returned is limited by the MaxPageSize constraint on the domain server. With paged results, there is no maximum. You can return all the results. You set the LDAP client page size in the **Search Options** dialog box. This value cannot exceed the MaxPageSize constraint imposed by the domain server.

If you use the **View Tree** option, LDAP does not use paged results and can only display objects limited by the MaxPageSize constraint in the results pane.

> **Note** The domain server imposes a limit on the maximum number of objects returned by a query in a single page. This is determined by the value of the MaxPageSize constraint in the LDAP Query Policy. The default value is 1000, which can be altered by changing the value in the default query policy or creating and assigning a new policy to the domain control. This limit means that any LDAP client can retrieve a maximum of 1,000 objects total if it does not use paged results, or a maximum of 1,000 objects per page if it uses paged results.

For more information about network protocols, see "Exchange 2000 Architecture" in this book.

API Components

Many of the features of Exchange 2000 are made possible by Windows application programming interface (API) components. These components are API function libraries, which are composed of Dynamic Link Libraries (DLLs). These DLLs contain the actual function calls that developers can use in their code. These functions can enhance the application while allowing the developer to reuse programming code and save time.

The API components used in Exchange are:

- ADO
- ADSI
- CDO
- OLEDB
- MAPI

ADO

ActiveX Data Objects (ADO) are used to navigate through Web Storage System. ADO allows the developer to access all Microsoft Database APIs, as well as search for and bind objects. ADO contains built-in objects that the developer can use to add features to the application. This permits developers to add dynamic objects to a static Web page.

Figure 27.1 shows the ADO object collection.

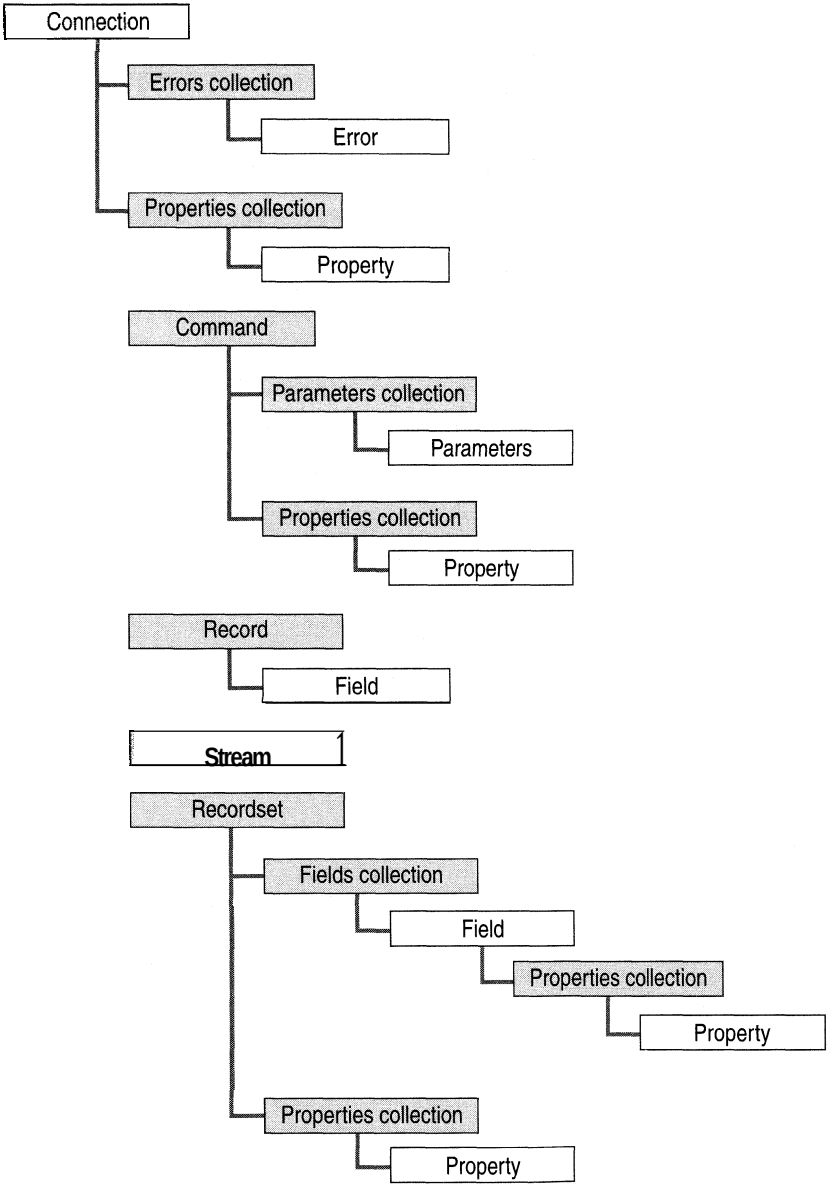


Figure 27.1 The ADO object collection

At the top of the ADO object collection diagram is the **Connection** object. It establishes and maintains a connection to a database or other data source. The open function of this object establishes the actual connection. This object manages transactions in applications. Beneath the **Connection** object is the **Error** object. It contains information on errors generated during an operation or command. A **Connection** object maintains a collection of **Error** objects.

The next object is the **Command** object. This provides commands in binary form.

Next is the **Parameter** object. This object enables the code to set or get a parameter to be used with a **Command** object. The **Command** object maintains a collection of **Parameter** objects.

Next is the **Record** object. This object opens an item or folder in Web Storage System using the item's URL. This object maintains a collection of fields.

The next object is the stream object. It accesses resources as bits of streaming data; when you open e-mail attachments, it streams the data instead of opening a large portion of data at once, thereby improving performance.

The next object in the diagram is the **Recordset** object. This object views one record at a time and provides functions for traversing a set of records. The object's open function provides access to a record set through a SQL statement, table name or **Command** object, and connection object.

Next is the **Field** object. It provides access to a **Recordset** object's field. Modifying the field value modifies the actual data in the data source. A **Recordset** object maintains a collection of **Field** objects.

The last object in the diagram is the **Property** object. It provides access to a property of the **Field, Command, Connection,** and **Recordset** objects. All of these objects maintain collections of **Property** objects. **Property** objects support dynamic properties, enabling a data source to add properties to an object at run-time.

Of all these objects, the **Connection** object is the most important. If an application slows down the network, it may be because the developer did not use the same ADO connection for multiple operations. Creating new connections throughout an application is both expensive and inefficient, especially when deploying ASP applications. The developer should create the **Connection** object at the beginning of the session. That object is then referred to throughout all of the subsequent ADO commands. By working with developers, you can bring up and resolve network resource issues in the early stages of application design.

Querying Web Storage System with ADO

To retrieve specific information from Web Storage System, you can use ADO to query only those records that you need. You must establish a connection explicitly with Web Storage System.

First, build the core SQL Select statement with the Web Storage System URL for the folder you want to query. Table 27.1 lists the SQL commands that are supported and unsupported in Web Storage System.

Table 27.1 Supported and unsupported SQL commands in Web Storage System

Supported SQL Commands	Unsupported SQL Commands
SELECT, SELECT *	JOIN
LIKE	MAX
WHERE	MIN
ORDER BY	SUM
GROUP BY	AVG
CONTAINS	
FREETEXT	
Column Aliasing	

Note The commands **CONTAINS** and **FREETEXT** work only if content indexing is turned on and your search item is indexed.

You need to specify the scope of the query in the **From** clause. **SCOPE** tells Web Storage System where to look for information. This can be either **DEEP** (this folder and all subfolders) or **SHALLOW** (this folder only) from the URL you specify. **SHALLOW** is the default.

The MAPI hierarchy (also called a MAPI tree) only supports **SHALLOW** traversals. You can perform deep traversals on non-MAPI public folders that you create.

The following is an example of an ADO query against Web Storage System:

- The **SELECT** statement takes two properties: the subject of a mail message and a custom property called **financial.**

  ```
  SELECT "urn:schemas:httpmail:subject",
  "urn:schemas:mydomain:financial"
  ```

- The **FROM** clause requires the keywords **Shallow** or **Deep** followed by **Traversal Of** and the URL of the folder.

  ```
  FROM SCOPE ( 'shallow traversal of file://backofficestorage/./public
  folders/accounts')
  ```

- Rather than return everything that matches, you can use the **WHERE** clause to restrict the returned records to those matching the subject.

  ```
  WHERE "urn:schemas:httpmail:subject" = 'loan'
  ```

When the query executes, the results are returned in a **Recordset** object. You can then perform whatever action you need to, such as **Copy,** or navigate through the contents.

Using ADO to Access Objects Through the ADSI Provider

The ADSI provider implements OLE DB interfaces that allow you to use ADO to access objects in compliant directories. The ADSI provider is read-only when used with ADO. You must create an ADO connection object and set its **Provider** property to **ADsDSOObject**. You can specify any string, including **""**, as the first argument connection string of the ADO connection object's **Open** function.

The connection object's **Execute** function's **CommandText** (first object) is an LDAP query composed of four elements separated by semicolons, in the following format:

```
<LDAP://server/adsidn>;ldapfilter;attributescsv;scope
```

Where:

- **sewer** is the name (or IP address) of the server hosting the directory.
- **adsidn** is the distinguished name of the starting point for your query expressed in ADsPath format with "/" separators and the root of the namespace to the left. You can also use an X-500 style attributed name format with the relative distinguished names separated by commas and the root of the name space to the right.
- **ldapfilter** is the LDAP filter string.
- **attributescsv** is **a** comma-separated list of names of the attributes to be returned for each row in the recordset.
- **scope** is either "base", "one level", or "subtrec".

To return the ADsPath, claw, and cn attributes of al! objects in the recipient containers on an Exchange server, you can use the following text in URL format:

```
LDAP:<//server/o=organization/ou=site/cn=recipients>;
    (objectClass=*);ADsPath,objectClass,cn;subtree"
```

or, in attributed-name format:

```
<LDAP://server/cn=recipients,ou=site,o=organization>, _
    (objectClass=*);ADsPath,objectClass;subtree
```

The following Microsoft Visual Basic sample code illustrates this query:

```
Dim conn As ADODB.Connection
    Dim rs As ADODB.Recordset
    Set conn = New ADODB.Connection
    conn.Provider = "ADSDSOObject"
    conn.Open "ADs Provider"
    Set rs = conn.Execute(
          "<LDAP://server/o=organization/ou=site/cn=recipients>; _
          (objectClass=*);ADsPath,objectClass,cn;subtree")
    While Not rs.EOF
       Debug.Print rs.Fields(0).Value, rs.Fields(1).Value, _
             rs.Fields(2).Value
       rs.MoveNext
    Wend
    conn.Close
```

ADSI

Active Directory Service Interface (ADSI) is used to add new users, modify attributes and browse Active Directory, as well as other activities. You can use ADSI to govern the lifecycle of abstract directory objects (such as computers and user groups) as well as access these objects, even if they are located on a different operating system. Using ADSI, you can update user information from a Web page and distribute it across a multi-operating system network.

ADSI uses the capabilities of directory services from different network providers to present a single set of directory service interfaces for accessing and managing network resources. You can use ADSI services to number and manage resources in a directory service, regardless of which network environment contains the resource. This can be an LDAP-based, NDS-based, or NTDS-based directory.

ADSI is a standard Windows-based interface for meta-directory applications. A meta-directory is a high-level network directory service designed to unite account and resource information from multiple network operating environments. The goal of ADSI is to synchronize this information across differing directories. Meta-directories extract security and account data from each directory, and then manage this information as part of an external database.

ADSI is comprised of a series of client-side DLLs that provide a common set of directory management functions. You can access these functions from almost any environment. As a result, you can add or remove user accounts, configure shared resources, and browse the directory tree all from a single, integrated console. Exchange 2000 uses ADSI to integrate application services with Active Directory's security and account functions. You can also access ADSI using Microsoft Visual Basic, Scripting Edition (VBScript) and Microsoft Windows Script Host to perform directory maintenance tasks.

ADSI allows application developers to query and update Active Directory. ADSI provides an abstraction of directory service contents, which makes it independent of the underlying directory service. The APIs and interfaces in ADSI conform to Component Object Model (COM) and support standard COM features, including automation.

ADSI allows different applications to communicate locally, which enables developers to create directory-enabled applications. ADSI is built on a provider-based model. While clients use the COM interfaces exposed by ADSI, providers implement the mapping between those COM interfaces and the underlying directory system. In this provider-based model, it can serve as the interface to a number of directory services, accessing each one through its own provider. This eliminates concerns about the underlying differences between various directory implementations or namespaces.

Because ADSI is a set of COM objects, you can use it with Visual Basic or Microsoft JScript to make Web-based applications directory-enabled. ADSI creates a directory structure in which you access any directory service using the same universal commands. This ADSI directory structure contains the following two basic types of objects:

- **Container** This object holds a number of other objects as either children or members.
- **Leaf** This object has neither children nor members.

The foundation object for this structure is the container named ADS://. This container stores a number of namespace containers. Each namespace represents a particular syntax that you use to gain access to a target directory service.

> **Note** The target service is not found in a namespace unless it has an ADSI provider installed or you write one of your own and give it a namespace.

Retrieving Information With ADSI

When you look for information, you must know whether the information is contained in the object instance or in the object's schema definition. The object name is contained in the object instance but a list of properties about the object is contained in the object's schema. The following example shows both an object instance and the object's schema definition:

```
Used Car Instance
Name: Horseless Carriage
Year Built: 1899
Engine: Steam
Used Car Object Instance
Attribute "Horseless Carriage", single value, data-type "text"
Attribute "Year Built", single value, data-type "Integer"
Attribute "Engine", single value, data-type "text"
```

In this example, the car's name, year built, and engine information are kept in the instance of a used car object. The schema contains definitions of the used car attributes, such as data-type and engine-type constraints.

In comparison, a SQL object definition (or object schema) is similar to a table definition and the object instance is similar to a particular row in the table. If the information crosses more than one table in SQL, the developer is responsible for managing those relationships. Using ADSI, the interface is responsible for the same relationships.

In ADSI, each object, regardless if it is an instance object or a schema object, has these properties: **Name, Parent, Schema, Class, GUID,** and **ADSPath.** These are all the properties you need to know to navigate through the hierarchical structure of the database. The **Name** property is the name of the object. The **Parent** property gives the location of the parent object. The **Schema** property governs the object. The **Class** property is the type of object. The class contains all the attributes that an object can have. The **GUID** property is the unique identifier for the object. The **ADSPath** is the location of the object you retrieve.

To get the property information for the default database, you must either know the name of the property or get the list of properties for the class. In the car example shown earlier, you have all the properties and their values. You can specify any object to have mandatory or optional properties. You must set all mandatory properties to create an object.

Using ADSI, you can create a temporary object outside of the source and set properties without affecting the performance of other requests to the data source. After you set the properties, you add the object to the database source with the **object.SetInfo** statement. This is also true for the **Update** command. Using ADSI, you retrieve the object once, set the properties, and then make comments about the updates with an **object.SetInfo** statement.

If you want to change or monitor a particular item in the directory structure, you need to bind an object to that item.

For example, if you want to change a user account, use the following syntax:

```
Set objUser = GetObject("WinNT://microsoft.com/Bob")
```

In this example, "microsoft.com" is the name of the local domain and "Bob" is the name of the account.

You can access the six core properties with the object/function syntax, and print them on your screen. When you are bound to a User object, you also have access to the properties and functions defined by the ADSI User object. The number of properties varies from one namespace provider to another.

It is important to know how to enumerate the objects in a container. The code varies slightly depending on whether the container objects are members or children of the container. For example, groups are children of a domain and users are members of a group.

For example, the code to enumerate the user "Executives" of the "GlobalHQ" group is as follows:

```
<HTML>
<BODY>
<%
  Dim oGroup
  Dim oExecutive
  Set oGroup = GetObject("WinNT://microsoft.com/GlobalHQ")
  For Each oExecutive in oGroup.Executive
    Response.Write oExecutive.Name & "<BR>"
  Next
%>
</BODY>
</HTML>
```

To list the children of an object, use a "For Each...Next" loop:

```
Set oContainer = GetObject("WinNT://Microsoft.com")
  For Each oChild in oContainer
    Response.Write oChild.Name & "<BR>"
  Next
```

These are the two important coding functions for any ADSI page. You can also write an ASP to number the members of the Mandatory Properties and Optional Properties of the schema for that object's class.

> **Note** ADSI occasionally generates an unexpected error message when sending a message. When there are no mandatory properties to number, you should use the **On Error Resume Next** property to avoid generating an error message and stopping the program.

For more information about what properties and functions each object class supports, see ADSI Help.

CDO

The Collaboration Data Object (CDO) library encapsulates difficult or repetitive operations, such as meeting requests, for manipulating items in Web Storage System. CDO is an API specification that defines objects, interfaces, functions, and properties for accessing messaging and collaboration data. It is also available in a series of object libraries. CDO for Windows 2000 is an SMTP-based library that provides base-messaging features. CDO comes with Microsoft Windows 2000 Professional and Windows 2000 Server. You can use these CDO objects in applications to save developers the time-consuming task of re-inventing commonly used components. You can use CDO for Windows 2000 to send e-mail messages and post discussion and news messages where no mailbox is required. It also provides extensive MIME content creation and management, and support for the SMTP and NNTP protocol stacks. You can use CDO for Windows 2000 for mail-intensive applications, as well as any application that needs to construct MIME content. CDO itself is not accessible remotely.

When Exchange 2000 installs, it adds an extensive set of CDO objects for Exchange object types and provides a larger set of objects than Windows 2000 CDO. Exchange 2000 CDO is compatible with all existing applications, including CDO 1.2x, MAPI, and Windows NT Server CDO. Exchange 2000 CDO provides objects for scheduling systems and resource booking, as well as support for e-mail, mailbox, voice-mail, fax, and pagers. You can also use Exchange 2000 CDO for server-based agents. CDO is an integral part of Microsoft's Rapid Application Development (RAD) strategy.

Exchange 2000 CDO works with ADO and OLE DB. You can use ADO and OLE DB for all generalized data access, such as rowsets, queries, hierarchy navigation, and access control lists (ACLs). You can use recordsets and rowsets in place of certain collections, such as Items, Folders, Recipients, Attachments, and Attendees.

You use CDO to access collaboration-specific data, such as Exchange-specific objects. CDO adopts Internet standards for content and protocols at the API level, including the following:

- RFC 822
- RFC 977
- RFC 1036
- MIME
- S/MIME
- HTML
- MHTML
- ICalendar (Internet standard for exchanging calendaring information)
- vCard
- SMTP
- NNTP
- LDAP
- HTTP

CDO provides a consistent programming model with the same objects and interfaces on the client and server side. For example, you can use the Outlook object model on the client side and CDO on the server side, and use the same application logic, behavior, and features across all applications.

Exchange 2000 CDO provides a convenient object model for managing folders, messages, Exchange mailboxes, appointments, contacts, and other items as well as the properties of those items. CDO is designed to operate in Web Storage System and supports Internet standards and protocols. It provides direct access to OLE DB interfaces for using Web Storage System.

CDO integrates with the ActiveX Data Objects (ADO) 2.5 component to provide you with a consistent data-access interface to Web Storage System and Active Directory. As you set values to CDO properties or access Web Storage System properties, CDO saves data to the correct Web Storage System or Active Directory locations.

Using IDataSource

CDO objects containing data have an IDataSource interface. This interface manages the data of a resource. The data is accessed as a property.

To open an existing item, use the following syntax:

```
prsExist.IDataSource.Open strURL, , _
   adModeReadWrite
```

To save a new item, use the following syntax:

```
prsNew.IDataSource.SaveTo strURL
```

Creating Contact Information

The CDO Person class represents contact information. CDO properties exist for most contact fields. You use the DataSource **SaveTo** or **SaveToContainer** function to save the contact. For example:

```
strURLFldFilms = conURLDomain & _
 "/public folders/Movie Production/microsoft/"
Set prsNew = New CDO.Person
With prsNew
    .FirstName = "Kim"
    .LastName = "Yoshida"
    .JobTitle = "Director"
    .HomeCity = "Los Angeles"
    .HomeState = "California"
    .Email = "kimy@microsoft.com"
    .DataSource.SaveToContainer
   strURLFldDoctors, , , adCreateOverwrite
End With
```

Scheduling

Scheduling uses Appointment, CalendarMessage, Attendee, DataSource, and Configuration. To perform scheduling in CDO, you must first configure for the organizer. You then set appointment properties, invite the attendees, save to a folder container and send a meeting request. For example:

```
Set appt = New CDO.Appointment
With appt
    .Configuration.Fields(cdoSendEmailAddress) = _
     "someone@microsoft.com"
    .Configuration.Fields.Update
    .StartTime = #10/9/1999 12:30:00 PM#
    .EndTime = #10/9/1999 1:30:00 PM#
```

```
.Subject = "Exchange 2000 Web Storage System - The Musical"
    .Location = "Film Theatre"
    Set att = .Attendees.Add
    att.Address = "bob@microsoft.com"
    att.Role = cdoRequiredParticipant
        Set msg = .CreateRequest
    msg.Message.Send
        .DataSource.SaveToContainer strURLContainer
End With
```

Sending Mail

You must have a sender and a recipient to send e-mail. First, configure the Sender, and then configure the Address with To, CC, or BCC properties. Separate multiple addresses with commas and set the subject and TextBody properties for the test message. Send the message by using the **Send** function. For example:

```
Set msg = New CDO.Message
With msg
    Configuration _
      .Fields(cdoSendEmailAddress) = "someone@microsoft.com"
    .Configuration.Fields.Update
    .To = "someone@microsoft.com"
    .Subject = "Movie Awards"
    .TextBody = "We've won an Oscar!"
    .Send
End With
```

Deciding When to Use ADO and CDO

You use ADO to navigate and manipulate records, set basic properties, and save generic items. You use Windows 2000 CDO for mathematical problems and repetitive operations.

OLE DB

The OLE DB API component provides a scriptable interface. Formerly known as Object Linking and Embedding Data Bases, OLE DB is a Microsoft technology that permits servers to share multimedia data, and allows applications and Microsoft software on those servers to exchange information and work together. Exchange 2000's OLE DB provider allows access to Web Storage System by using ADO.

MAPI

Extended MAPI version 1.0 is an industry-wide standard **for** writing messages, e-mail, and workflow applications. Exchange 2000 and many other common messaging programs use MAPI. MAPI interfaces create and access diverse messaging applications and messaging systems, offering a uniform environment for development and use, and providing true independence for both.

You can use simple MAPI if you can call an entry point in a DLL and pass a data record conforming to an externally-defined structure. You can use full MAPI if you can implement or manipulate arbitrary objects conforming to the Microsoft Component Object Model (COM).

Formats

Formats specify the properties (particularly visible properties) of an object. For example, word processing applications allow you to format text, which involves specifying the font, alignment, margins, and other properties. Exchange 2000 uses the following formats:

- RFC 822
- MIME
- HTML
- XML

RFC 822

RFC 822 is a message representation protocol which specifies details about message headers, but which leaves the message content, or message body, as ASCII text. In the TCP/IP protocol suite, e-mail is in RFC 822 format.

While SMTP has remained unchanged over the years, the Internet community has made several changes in the way it uses SMTP. In particular, the conversion to Domain Name System (DNS) has caused changes in address formats and in mail routing. RFC 822 specifies the Internet standard format for e-mail messages. RFC 822 replaces RFC 733, which is a previous standard that some organizations might use. Occasionally, the two formats are referred to by their numbers 822 or 733.

Some non-Internet mail environments with mail transfer protocols other than SMTP use RFC 822. SMTP is adapted for use in some non-Internet environments.

MIME

Multipurpose Internet Mail Extensions (MIME) is the standard protocol for including non-text information in e-mail, such as binary data, audio data, video data, and foreign language text that cannot be represented in ASCII text. MIME supports the transmission of mixed-media messages across TCP/IP networks.

HTML

HTML is the authoring language used to create documents on the World Wide Web. HTML defines the structure and layout of a Web document by using a variety of tags and attributes.

XML

Extensible Markup Language (XML) receives and responds to Internet data requests. XML is a foundation for building applications. XML files are easy to create in Web Storage System. The XML Data Object Model can save to a Web Storage System URL using the object's **Save** function. You can create XML **Data Object Model** objects based on Web Storage System properties and end-user choices from a simple HTML form. You can also store XML as the value for a Web Storage System property.

Extending Exchange data with custom data is simple. You can combine it with e-mail, calendaring (a protocol used with CDO), tasks, and contacts. You can use XML to add important metadata to documents, increasing their usability and manageability. You can also use XML to define data structurally in a platform-independent, language-independent manner.

By convention, property names for Web Storage System items are prefixed with a Uniform Resource Identifier (URI). These URIs are a set of unique names used to ensure that Web Storage System renders globally-unique names. Item properties are requested and returned in WebDAV sessions as XML elements that use URI namespaces. A URI for declaring namespaces can either be a URL or a Uniform Resource Name (URN).

For searches from Web Storage System clients, a SQL query can be sent via an Internet request by using the **XMLHttpRequest** object. The XML parser imports the document into a dynamic data structure that can be recursively traversed. The HTTP 1.1 **Search** function sends the SQL query as the value of an XML node object. You can then manipulate a returned **XML Data Object Model** object with the search results.

Microsoft Web Storage System

The Web Storage System is a semi-structured, hierarchical database that stores and manages documents, e-mail messages, Web pages, multimedia streams, data elements, and other items. Developers can access Web Storage System using any programming language that supports COM and the Exchange protocols listed earlier in this chapter.

Similar to file system directories and files, Web Storage System is a hierarchy of folders. Each folder in Web Storage System is a collection of items, including other folders. These folders and items are accessible by URLs.

With Web Storage System. you can store an item and maintain property information about that item, including controlling its security and display. These properties are actual fields on the item's row in Web Storage System. An item's body (if present), whether text or binary, is itself a field on the row and is accessible as a stream. You can use ADOs to move through Web Storage System hierarchy and obtain data using **Record** and **Recordset** objects. You use ADO or OLE DB to find and obtain objects from Web Storage System items, as well as create, delete, move, and copy folders and items. You can also manipulate data using HTTPDAV, which allows the user to add, copy, modify, move and search for items.

Web Storage System Schema

The Web Storage System defines a schema for the formats described earlier in this chapter. You can define a custom set of properties for your items and control their use. The **content class** property indicates the intent or purpose of the item. Therefore, a property does not need previously defined schema before it can be added to Web Storage System. The Web Storage System schema provides the means for Web applications and search queries to efficiently access your data.

IIS Integration

Exchange 2000 fully integrates with the Internet because of the protocol support in IIS. There is a specialized high-speed link between IIS and Exchange. IIS determines which virtual path to use when a request is received from the network. For example, if there is a request for http:Nserver/exchange, the Exchange virtual path route's Internet Server Application Programming Interface (ISAPI) register (DAVX.dll) responds to the request. That DLL in turn can communicate over a dedicated high-speed channel to an HTTP address that resides in Web Storage System process. The HTTP address takes the request and looks into Web Storage System. Depending on the nature of the request, the address returns a row set, delete, or move to an item, and then returns a response to the ISAPI register. The register then takes the results of the request, renders the data, and either returns the result back to IIS or returns a response code to the user.

You can direct IIS to gain access to Web Storage System using the Win32 file system, but this only allows the user to add or delete files. Figure 27.2 illustrates the network request process using a loan approval ASP page as an example.

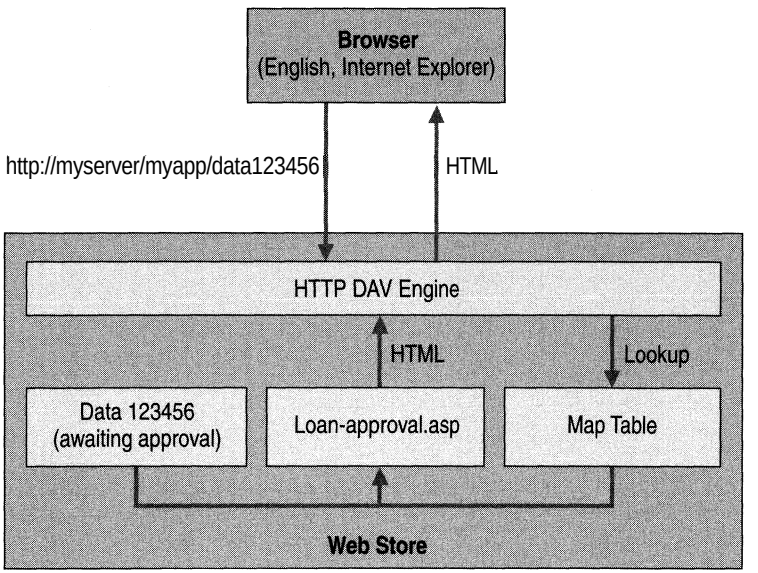


Figure 27.2 Network request process flow of a loan approval ASP page

The Exchange Server Active Directory Schema

Active Directory provides you with a way of extending the schema. This differs from the Exchange Server 5.5 Schema, which does not allow you to add new classes or attributes. To write applications that access Active Directory, it is useful to have an understanding of the Active Directory schema. The directory schema is a collection of descriptions, or rules, that define characteristics of objects in the directory. The schema defines available object classes in the directory, relationships between object classes, attributes of each object class, and specific characteristics of attributes and classes. For more information about Active Directory Services, go to http://www.microsoft.com/exchange.

Exchange 2000 Web-Based Collaboration Applications

Exchange 2000 enables developers to create Web-based collaboration applications. These applications allow developers to make full use of Exchange 2000 by using Internet technology.

The Microsoft Web-based Collaboration Applications are as follows:

- Microsoft Windows 2000 Server

  Microsoft Active Server Pages (ASP)

  Microsoft Internet Explorer
- Microsoft Internet Information Services (IIS)
- Microsoft Outlook 2000

These technologies are unified through Windows 2000 Active Directory, and give developers the ability to rapidly develop Web-based applications.

Microsoft Windows 2000 Server

Exchange 2000 is integrated with Microsoft Windows 2000 Server to utilize its security features, stability, and scalability. Windows 2000 Server and Exchange 2000 Server come with CDO. Exchange 2000 Server comes with its own CDO library, which provides a superset for the CDO library in Windows 2000.

Microsoft Internet Information Services

Microsoft Internet Information Services (IIS) simplifies communications and delivering Web sites to Web browsers. IIS is the only Web server integrated into Windows 2000 Server. IIS is the front-end for Internet protocols in Exchange 2000. Protocol plug-ins in IIS communicate with Exchange 2000 Server about incoming and outgoing e-mail. This integration is possible by using the Exchange 2000 Installable File System (ExIFS), which allows large amounts of data to pass quickly between IIS and Exchange.

IIS supports ISAPI, which enables programmers to develop Web-based applications that run much faster than conventional Common Gateway Interface (CGI) programs because they are more tightly integrated with the Web server. ISAPI extends Internet services, such as IIS, by using extensions and filters. Using an ISAPI extension, a browser can send information entered in a form and the extension will return a complete HTML page built from code. In an ISAPI filter, all browser data. both inbound and outbound, can be modified before or after the server processes it. Extension and filter DLLs must conform to a specification. For the DLL to work there are specific functions that must be implemented.

In previous versions of Exchange, the information store manages the databases and client access protocols such as IMAP4, POP3, MAPI, and NNTP. In Exchange 2000, the Internet access protocols are removed from the store and are instead managed by IIS. Deploying a front-end/back-end configuration makes it possible to manage the Internet access protocols on a server that is separate from the one on which the store and databases run. Essentially, a bank of protocol servers process the incoming client connections while the store servers are dedicated to running the databases. The main benefits of a front-endhack-end configuration are a unified namespace and reduced network bandwidth traffic for SSL encryption.

Active Server Pages (ASPs)

An Active Server Page (ASP) is an HTML page with server-executed script embedded in the page. ASP is an open architecture application that does not need to be compiled. This makes it possible for users to combine HTML, scripts, and reusable ActiveX server components to create Web-based collaboration applications. ASP enables server scripting for IIS with native support for both VBScript and Microsoft JScript.

When IIS receives a request for a page with an .asp extension, script within the HTML page is interpreted and presented back to the browser as HTML. Script is designated by <SCRIPT> and <%..%> tags. Any scripting language can be used as the script language, but the most common languages are either VBScript or JScript. Scripts are not stored in compiled form. They are interpreted when the .asp file is requested from the server.

An ASP page can run under the context of the user requesting the page, the IIS_USER account or another account if configured. A common scenario is to require Windows 2000 authentication for the page so that it runs when the user makes the request. In this model, the user running the page must have the access control permissions to do the operation.

The following is a sample ASP page using ADO to search Web Storage System:

```
<%@ Language=VBScript %>
<HTML>
<HEAD>
<META NAME="GENERATOR" Content="Microsoft Visual Studio 6.0">
</HEAD>
<BODY>
<%
Dim rec 'As ADODB.Record
Dim rst 'As ADODB.Recordset
Dim urlFolder 'As String
```

```
'Build the URL
urlFolder = _
"file://./backofficestorage/nwtraders/microsoft.com/Management/"

'Create a Record object
Set rec = server.createobject("ADODB.Record")

'Open the Management application folder
rec.Open urlFolder

'Extract the contents of the folder
Set rst = rec.GetChildren
%>

<h1>Folder List:</h1>
<hr>

<%
' Read the URL for each child record:
Do Until rst.EOF %>
   <%=rst.Fields("DAV:href")%><p>
   <%rst.MoveNext
Loop

'Close the objects and release memory
rst.Close
rec.Close
Set rst = Nothing
Set rec = Nothing
%>

</BODY>
</HTML>
```

Microsoft Internet Explorer

Microsoft Internet Explorer version *5.0* is a key component of Web-based application development in Exchange 2000. You can access Web Storage System in Exchange using a Web browser such as Internet Explorer. By providing a URL, you can browse from the Internet straight to an item.

The following are examples of items to which you can browse:

- Private mailbox folder:
  `http://domainname/exchange/yourname/inbox`
- Public store:
  `http://domainname/public/someFolder`
- Mailbox item:
  `http://domainname/exchange/yourname/inbox/hello.eml`

Always type the domain name in a front-endhack-end server configuration.

The Web Storage System renders an HTML 3.2 view of a folder's items. You can enhance these views in a user interface using Web Storage System forms.

Internet Explorer *5.0* includes the *XML* COM component, a component containing an XML parser, and related tools. Internet Explorer 5 has an object model that works with XML to communicate with WebDAV servers.

SMTP and NNTP Inbound and Outbound Events

SMTP and NNTP inbound and outbound events refer to information that is destined for or leaving Exchange, respectively. With Inbound Events, IIS creates a new file in ExIFS and writes to it. ExIFS returns a list of pages to which the message was written. The list of pages is passed to the Exchange Server process where the Extensible Server Engine (ESE) commits the pages by logging the information, and page checksums (used to check data integrity) are stored in the EDB file.

In Outbound Events, the Exchange server receives the list of pages referenced by an outgoing message and passes the list of pages to IIS. In turn, IIS opens a file in ExIFS using the list of pages. The message can now be transmitted quickly.

Note Page checksums are not verified during transmission.

Space Allocation

Space allocation refers to making space available in the database. The ESE tracks pages in the streaming file that are currently committed. These pages are reserved in ExIFS. It allocates space for new messages from its reserved space and requests space from ESE when necessary.

File Semantics

File semantics refers to the wording used to call files. ExIFS can place Win32 file calls into the database. A call such as **FindFirstFile** can look in a mailbox folder for a list of messages. ExIFS can display the contents of a database as ordinary files.

> **Note** In IIS version 4 and IIS version 5, the metabase stores configuration information. You gain access to the metabase by using ADSI (including Script). All SMTP configurations are stored in the metabase.

Microsoft Outlook 2000

Outlook 2000 provides tools for creating programming and non-programming collaboration applications. The Desktop Information Manager included in Outlook 2000 centralizes all communication, organization, and management of information. Outlook 2000 supports multiple e-mail providers, scheduling, task management, journals, contacts, and the creation of personal folders for storing and categorizing information.

For application developers, the Outlook 2000 application design environment provides advanced design features and tools. Some of these features include:

32-bit Forms

You can display and update forms quickly because Outlook 2000 forms are 32-bit and form definitions are usually less than 12 kilobytes (KB).

Instant Switching Between Forms and Run Time

Designers can switch instantly between form design and run time so they can quickly modify and test forms.

Advanced Fields and Views

Fields in an Outlook 2000 form or view can include calculated expressions, validation formulas, and number formatting.

Built-in Forms Modules

Forms modules support all the features of the built-in Outlook 2000 forms, including digital signatures and encryption. You do not lose standard user interface design and features when using a custom form. You can send custom forms from outside the organization. Users can embed form definitions or layouts in a message so they can send them inside or outside their organization to others who might or might not have the form definition installed.

Extendable Forms

Developers can extend forms with existing Microsoft ActiveX or OLE controls and with the VBScript language.

The Outlook 2000 design environment enhances the design features offered with Microsoft Exchange Client. Outlook 2000 provides tools to create and design folders and customize forms. These tools include Outlook Forms Designer, VBScript, and the Microsoft Visual Basic Expression Service.

Outlook Forms Designer

Outlook Forms Designer is the integrated development environment included in Outlook 2000. You can access it in any Outlook 2000 form by clicking **Form** on the **Tools** menu. You can publish forms created using the Outlook Forms Designer in folders and save them as .oft files.

Visual Basic Expression Service

Visual Basic Expression Service is built in to the Outlook Forms Designer environment, allowing designers to create validated, formula and combination fields. For example, using Visual Basic Expression Service, a designer can create a formula field for a timecard form that calculates the total number of hours worked in a week.

Instant Collaboration

Instant Collaborations are built-in application modules included in Outlook 2000. You can customize these modules. To create an instant collaboration, you must place a module in an Exchange public folder. The built-in modules are: Calendar, Tasks, Journal, and Contacts. Built-in modules contain predetermined function and design settings, and they require only the services of a public folder to become an instant collaboration.

To create an instant collaboration, you must first create or locate a public folder to contain the module. You then copy a built-in module to the public folder. The combination of the Outlook 2000 module and Exchange public folder results in a collaboration application.

You can also create collaboration applications using the built-in modules included in Outlook 2000. For example, to create an application to track individuals or companies, you can use the Contacts module to take advantage of the built-in views and fields.

Built-in applications have a variety of predefined views that you can easily integrate with the Journal module. Contacts applications offer e-mail and Web integration. For example, users can click on a hyperlink to go to a site on the Internet.

A contacts application provides a group of users with a shared database that stores information about clients, friends, or vendors. To make it a collaboration, users can post to it in a public folder.

Creating an application using one of the built-in modules does not require a developer's expertise. The Outlook Forms Designer and VBScript are available for more complex applications.

Calendar Module

When you copy or create the Calendar Module in a public folder, users can share, post, and update schedules for activities such as training classes, sporting events and company functions. For example, product launch milestones such as trade shows, product ship dates, or press tours can be posted to a Calendar folder for group viewing.

Tasks Module

When you copy or create the Tasks Module in a public folder, team members can share a common task list that displays who is responsible for a task and the status of the task. For example, a project manager can create a public Tasks folder that team members can update when tasks are completed or significant progress is made. This provides the product manager with up-to-date information on the status of a project.

Contacts Module

When you copy or create the Contacts Module in a public folder, users can add to, update, and share a list of contacts. For example, the sales department can share a list of leads or the entire company can share a list of vendor contacts.

Journal Module

When you copy or create the Journal Module in a public folder, users can log and track information such as the amount of time an individual spends on a particular task, on a project, or with a specific customer. You can set the Journal to log and store Office documents, contact calls, e-mail, and other communications.

Notes Module

You can share the Notes module or use it privately. It is the graphical interface equivalent to self-stick notes. As with the other built-in modules, copying or creating this module in a public folder allows users shared access. The color and category can be customized for easy retrieval, and the note can be forwarded as a message.

Event Support in Exchange 2000

r

Events allow applications to define custom actions on particula events. For example, a workflow application can initiate a workflow process when an object is saved to a particular folder.

There are three types of events: Protocol events, Transport events, and Web Storage System events. Protocol events are actions or occurrences that take place at the protocol level. Transport events extend the message processing system. Web Storage System events are OLE DB–based server events that provide an API through which developers can extend the features of Web Storage System. Web Storage System events include synchronous, asynchronous and system events. Synchronous events occur before the change is committed to the store. Asynchronous events are essentially notifications that happen after the item is saved or deleted. There are three Web Storage System events: starting Web Storage System, stopping Web Storage System, and the timer.

You can control which users can register event sinks and what the event sinks can access. For more information about Event Security, please go to http://www.microsoft.com/exchange.

Using Event Sinks in Exchange 2000

Exchange 2000 enables the user to write code to gain access to SMTP and NNTP stacks. For example, you might want to add a disclaimer to all out-going e-mail. With Exchange 2000, you can write the same CDO sink in any COM-compatible language. This allows you to create a mail system on top of the Windows 2000 SMTP service, as well as extending Exchange 2000 and Windows 2000 SMTP services.

You can build event sinks by creating DLLs in either Microsoft Visual C++ or Visual Basic. You can also register event sinks using simple script.

Web Storage System Events

Web Storage System events provide a mechanism through which applications can execute code whenever an item is saved, deleted, moved, copied or modified. Web Storage System events run on the server of the client that causes the event to happen. By using Web Storage System events, an application developer can build application logic that executes if the client is a MAPI client such as Outlook, an HTTP client such as Internet Explorer or Outlook Web Access, or a Win32 client such as any Microsoft Office application. Developers can build application logic that executes if the item is delivered using SMTP, TMAP or NNTP.

Separating the application logic from the client accessing the data and moving it onto the server relieves the developer from the task of rewriting the code for every client that can access the data. It also ensures that the application logic runs whether or not the client is connected when the item arrives in the Mailbox Store or the client has the application installed.

The Web Storage System causes a number of different events. These events are either synchronous, asynchronous or system events. Table 27.2 shows Web Storage System events supported by Exchange 2000.

Table 27.2 Web Storage System events supported by Exchange 2000

Asynchronous Events	Synchronous Events	System Events
OnSave	OnSyncSave	OnTimer
OnDelete	OnSyncDelete	MDBOnStartUp
		MDBOnShutDown

The Web Storage System guarantees that asynchronous events are delivered even if the store is shut down before the event fires. When this occurs, the event fires when the store restarts.

Developers can write code in Visual Basic, Visual C++, or any other language that supports COM.

Synchronous Events

An application that processes a synchronous event can modify the item before it is saved to the store, prevented from being saved to the store, or deleted from the store. Because the application is called before the item is saved, the application can modify the item before it is perceived or accessed by the client.

Synchronous events occur within the context of a local transaction. As such, they are called once before and once after the transaction is committed. The first time an event sink is called is referred to as the *begin phase*. During the begin phase, the event sink has full read/write access to the item that caused the event. The second time the event sink is called is referred to as the commit or *abort phase*. During this phase the item is read-only.

The Web Storage System provides two synchronous events: synchronous save and synchronous delete. Synchronous save is called whenever an item is saved or delivered into the store. An application that attempts to process synchronous save events must implement an event sink with the following signature:

```
Private Sub IExStoreSyncEvents_OnSyncSave(
ByVal pEventInfo As Exoledb.IExStoreEventInfo,
ByVal bstrURLItem As String,
ByVal lFlags As Long
)
' Your code here
End Sub
```

The first parameter contains the event information structure. The second parameter is a string that contains the URL of the item on which the event is occurring, and the third parameter is a long data type containing a series of bits that describe the type of save event and the phase of the event.

This provides functions and properties through which to abort the transaction, get an interface with the item being saved, get an interface with the registration item, and get an interface with the session.

To get an ADO record with the item causing the event specified in the pEventInfo parameter, you must first get a dispatch interface with the event information structure as follows:

```
Dim DispInfo As IExStoreDispEventInfo
Set DispInfo = pEventInfo
```

Then you simply get the event item as a record:

```
Dim ado-rec As ADODB.Record
Set ado-rec = DispInfo.EventRecord
```

When you have a record interface with the event item you can read, write or modify any property of the item. Because you are modifying the item within a synchronous save event, all changes that you make to the item occur before the item is saved to the store.

```
Dim subject As String
subject = ado_rec.Fields("urn:schemas:mailheader:subject")
ado_rec.Fields("urn:schemas:mailheader:subject") = "This is a demo(" &
subject &")"
ado_rec.Fields.Update
```

It is important to understand that new items you save to Web Storage System do not exist until the transaction is committed. Because the item does not exist, the URL to the item is invalid until it is saved. The following code demonstrates this and shows how to use the flags passed to the event sink to determine what caused the event and the item's current state.

```
Private Sub IExStoreSyncEvents_OnSyncSave1(
ByVal pEventInfo As Exoledb.IExStoreEventInfo, _
ByVal bstrURLItem As String,
ByVal lFlags As Long)

    If (lFlags And Exoledb.EVT_SYNC_BEGIN) > 0 Then
        If (lFlags And Exoledb.EVT_NEW_ITEM) > 0 Then
            'This is a new item so bstrURLItem is NULL
        Else
            'This is not a new item so
            'bstrURLItem contains the URL to the item
        End If
    Else
        If (lFlags And Exoledb.EVT_SYNC_COMMITTED) > 0 Then
            'The item has been saved
            'so bstrURLItem contains the URL to the item
```

```
        Else
            If (lFlags And Exoledb.EVT_SYNC_ABORTED) > 0 Then
                'the item will not be saved
            End If
        End If
    End If

End Sub
```

The flags parameter is set to Exoledb.EVT_SYNC_BEGIN the first time the event sink is called. The Exoledb.EVT_NEW_ITEM is also set the first time the item is saved, and the URL to the item is NULL. On subsequent saves, the Exoledb.EVT_NEW_ITEM is not set and the URL contains a valid URL to the modified item.

The second time the event sink is called, the flags parameter is either set to Exoledb.EVT_SYNC_COMMITTED if the item is committed to the store or Exoledb.EVT_SYNC_ABORTED if the transaction is aborted.

Synchronous delete is called whenever an item is deleted from the store. An application that attempts to process synchronous delete events must implement an event sink with the following signature:

```
Private Sub IExStoreSyncEvents_OnSyncDelete(ByVal pEventInfo As _
Exoledb.IExStoreEventInfo, ByVal bstrURLItem As String, ByVal lFlags As
Long)
' Your code here
End Sub
```

Because an item must exist before it can be deleted, the URL passed to the event sink always points to a valid item in the store. You can abort the transaction by calling **AbortChange** during the begin phase of a synchronous event sink. To conditionally abort a delete based on a property of the item, you can check the value of the property and then conditionally abort the transaction. For example:

```
Dim DispInfo As IExStoreDispEventInfo
Dim ado-rec As ADODB.Record
Set ado-rec = DispInfo.EventRecord
Set DispInfo = pEventInfo
Dim subject As String
If (InStr( ado_rec.Fields("http://schemas.mycompany.com/myprop"), "ready
to delete") > 0 Then
DispInfo.AbortChange 1
End If
```

This example aborts the transaction and returns the error code passed to the **AbortChange** function. This is particularly useful when controlling the characteristics of a document that is involved in a workflow process. A typical workflow can use ACLs to determine both which users can read or write to a document and the status of the document. One or more workflow properties on the document (or even on another document) record this information to control which users can delete the document.

Asynchronous Events

Asynchronous events happen after the item is saved to the store. They serve as a notification when it is not necessary to modify the item before it is saved. It is important to ensure that your code does not delay the action more than necessary. Code that executes within a synchronous event must complete before the item can be saved to the store. Therefore, applications typically defer actions that do not directly affect the item to an asynchronous event. For example, you can change the workflow state of a document in a synchronous event, but you should defer updating a Web page that advertises the existence and state of the document to an asynchronous event. You should not send out notification e-mail until the item is saved and the asynchronous event fires.

Web Storage System provides two asynchronous events: asynchronous save and asynchronous delete. Asynchronous save is called whenever an item is saved or delivered to the store. An application that attempts to process asynchronous save events must implement an event sink with the following signature:

```
Private Sub IExStoreAsyncEvents_OnSave(
ByVal pEventInfo As IExStoreEventInfo,
ByVal bstrURLItem As String,
ByVal lFlags As Long)
' Your code here
End Sub
```

The first parameter contains the event information structure. This provides an interface to the registration. It does not provide an interface to the event item. To access the item you must bind to it using the URL passed in the second parameter by using an ADO record:

```
Dim ado-rec As New ADODB.Record
ado_rec.Open bstrURLItem, , adModeRead, adFailIfNotExists
```

Asynchronous delete is called whenever an item is deleted from the store. An application that attempts to process asynchronous delete events must implement an event sink with the following signature:

```
Private Sub IExStoreAsyncEvents_OnDelete(
ByVal pEventInfo As Exoledb.IExStoreEventInfo,
ByVal bstrURLItem As String,
ByVal lFlags As Long)
' Your code here
End Sub
```

Registering for Web Storage System Events

When you write and install an event sink on Web Storage System server you must then register to be called whenever a specific event occurs. To create an event registration, save the item with a content class of **urn:content-class:storeeventreg** to the folder in which you want the event to fire using ADO.

To create an event registration:

```
Dim rBinding
Set rBinding = CreateObject("ADODB.Record")
rBinding.Open strBindingURL, cSess, 3, 33554432
```

To set the content class to make this a registration item:

```
rBinding.Fields("DAV:contentclass") = "urn:content-class:storeeventreg"
```

To set the name of the event to catch:

```
rBinding.Fields("http://schemas.microsoft.com/exchange/events/EventMethod
") = "OnSyncSave"
```

To set the program ID of the sink that services the event:

```
rBinding.Fields("http://schemas.microsoft.com/exchange/events/SinkClass")
= "MySink.Sink"
```

To enable the event registration and save the registration item to the store:

```
rBinding.Fields("http://schemas.microsoft.com/exchange/events/Enabled") =
bEnabled
rBinding.Fields.Update
rBinding.Close
Set rBinding = Nothing
```

Accessing the Item

The item that fires a synchronous event is returned to the application when the event fires. You can access this item using ADO or OLE DB.

To access a message attachment:

```
Dim cdo-attachment As CDO.IBodyPart
Dim cdo_msg As New CDO.Message
cdo_msg.DataSource.OpenObject ado-rec, cdoAdoRecord
For Each cdo-attachment In cdo_msg.Attachments
Dim attachment-stm As New Stream
```

To create the URL for the attachment:

```
Dim dir As String
Dim newItemURL As String
```

```
Dim filename As String
dir = ado_rec.Fields("DAV:parentname")
filename =
cdo_attachment.Fields("urn:schemas:httpmail:attachmentfilename")
newItemURL = dir & "/" & filename
```

To create a new item in the store:

```
Dim newItem_rec As New ADODB.Record
Dim newItem_stm As New Stream
newItem_rec.Open newItemURL, , adModeReadWrite, adCreateNonCollection
Or
adCreateOverwrite
```

To open a stream onto the new item:

```
newItem_stm.Open newItem_rec, adModeReadWrite, adOpenStreamFromRecord
```

Get a stream onto the attachment:

```
Set attachment-stm = cdo_attachment.GetDecodedContentStream
```

To copy the attachment to the new item:

```
attachment_stm.CopyTo newItem_stm
```

To save all:

```
newItem_stm.Flush
attachment_stm.Flush
```

To clean up:

```
Set newItem_rec = Nothing
Set newItem_stm = Nothing
Set attachment-stm = Nothing
Next cdo_attachment
Set cdo_msg = Nothing
```

SMTP Transport Events

SMTP Transport Events extend the message processing system (address resolution, restrictions processing, and message routing). The message is received in its entirety into the SMTP transport and is handed off to the customized code. The code runs on each message, providing read/write access to the envelope and body, and executes in process with IIS. You can use transport events to pass an incoming or outgoing message through a custom process before the message is stored or relayed to a recipient. Transport events can fundamentally alter the structure of the message by adding additional information such as disclaimers, or by passing the message to a compression agent before submitting information to the next destination. Table 27.3 shows the commonly used transport events in Microsoft Exchange 2000.

Table 27.3 Commonly used SMTP transport events

Event	Description
OnMessageSubmission	This event fires for every message that is submitted to the system. This occurs whether the message originated from an SMTP client, a MAPI client, or another client. In this event, you can write message-specific events such as virus checking and deleting mail from specific recipients.
OnSyncMessagePreCategorize	This event fires after the submission of a message. The only difference between this event and the OnMessageSubmission event is that no recipients should be added to a message in this event.
OnMessagePostCategorize	This event fires after the categorization of each message. Once this event fires, the categorized properties of a message are available to sinks. For example, if a sink needs to examine each recipient for a message, it should hook onto this event.

SMTP Protocol Events

You use protocol events for billing and charge-back, based on the number of connections and the time of connections. These events are also useful for monitoring the system using SMTP and for implementing new SMTP commands. Protocol events extend the SMTP protocol and modify existing SMTP protocol commands and responses.

SMTP Transport Components

SMTP services contain the following transport components:

- **Store Driver** Provides backing store for MailMsg objects passing through the service, and delivers local messages
- **Advanced Queue** Provides the queue management and logic for message delivery, routing, and relay
- **Categorizer** Provides categorization services for inbound messages, such as distribution list expansion using the LDAP and Active Directory

  Router Provides routing logic for outbound message relay, such as the Next Hop to which the message should be relayed

For more information about implementing these components, go to http://www.microsoft.com/exchange.

Event Bindings

Bindings are stored in the metabase as /smtpsvc/<instance number>/EventManager. You use ADSI to modify the metabase to create bindings. These are tools for registration, de-registration, debugging, and list binding. The Event dispatcher uses bindings to match sinks to events. The event rules filter the messages to determine which sink is applicable, if any. It then determines the priority for each event. Event sinks are executed in order of priority. Events with the same level of priority are executed in random order.

SMTP Protocol Event Sinks

You can extend SMTP features by using protocol sinks. For example, you can add a protocol keyword, reject mail from specific recipients, and exchange information with other SMTP event sinks on the other end of an SMTP connection. You can also restrict the size of messages that can be sent in your organization. Each protocol extension is a collection of sinks that perform aspects of the intended protocol command.

A sink can implement one or more sink interfaces. It fires depending on the session category. Inbound and Outbound categories include On Inbound Command, On Server Response, Session Start, Message Start, Per-Recipient, Before Data, and Session End. Protocol event sinks include:

- **IsmtpInCommandSink** This is called when inbound an SMTP command is received.
- **IsmtpOutCommandSink** This is called when outbound an SMTP command is sent.
- **IsmtpServerResponseSink** This is called when a server responds to a previously sent command.

SMTP and NNTP Transport Event Sinks with CDO

SMTP is an application-level TCP/IP protocol that supports text-oriented e-mail between computers supporting Message Handling Service. NNTP is the standard for Internet exchange of USENET messages. These services provide customized synchronous transport event sinks. A transport event occurs when message data is transported into or out of the services. All messages go through events, including X.400 and RPC inbound or outbound. One common example of an inbound transport event is the arrival of a message to the SMTP service either over the network or in the SMTP pickup directory. When this transport event occurs, a source for this event notifies transport event sinks, each of which can take some type of action based on the data content of the message. The system performs no other action with the message data after passing it to the sink until the sink completes and returns. While the sink is executing, it has exclusive control of the message data and state.

The architecture for the SMTP and NNTP source and sink event system is based on the Component Object Model (COM). A transport event sink is any COM object that implements the appropriate COM interface for the event. An event *source* is a process or thread that notifies sinks that are using the COM run time when a particular event occurs. The source needs no information about the internal details of the sink; it only requires that the sink be a COM object and that it returns the correct COM interface. Table 27.4 shows a number of transport event sinks.

Table 27.4 Transport event sinks

Event	Interface	Description
OnArrival	**ISMTPOnArrival**	Processes the arrival of messages to the SMTP service
OnPostEarly	**INNTPOnPostEarly**	Processes the initial arrival of newsgroup headers for messages to the NNTP service
OnPost	**INNTPOnPost**	Processes the arrival of new posts to the NNTP service
OnPostFinal	**INNTPOnPostFinal**	Processes all new posts that have been committed to disk or stored by the NNTP service

Event Binding Rules

You should limit protocol events to the command that you want to "hook"—that is, either send or respond to:

- To hook MAIL FROM:, hook "MAIL"
- To hook RCPT TO:, hook "RCPT

Transport events limit the event call to certain messages. Transport events are case-insensitive and can filter on **MAIL FROM=, RCPT TO=,** and **EHLO=**. For example:

```
MAIL FROM= user1@domain; RCPT TO= user2@domain
```

> **Note** These messages cannot be submitted to a Web Storage System that has a filter. You use rules only for SMTP messages.

In this example, the sink fires only on messages from user1@domain or to user2@domain. It is possible to substitute wildcards in this example (*@domain, user@random*) or have multiple RCPT TO:s, multiple MAIL FROM:s, and multiple EHLOs. In each case, you choose between FROM or TO.

CDO provides transport event scripting hosts that enable you to implement sinks using a scripting language. The process you use is as follows:

1. Write the script code that implements the appropriate event. Just as one sink COM class can implement multiple event functions, one script file can contain code for multiple (or all) events.
2. Register the appropriate CDO event scripting host as the event sink. The scripting host in turn dispatches the event function invocation to the scripting engine that runs your script code.
3. In **Sink Binding Properties,** set the **ScriptName** property so that the binding points to the location of the file containing the script you wrote in the first step. The CDO scripting host will execute the appropriate subroutine or function contained in the script when the event occurs.

Figure 27.3 illustrates the event sink pathway in Web Storage System.

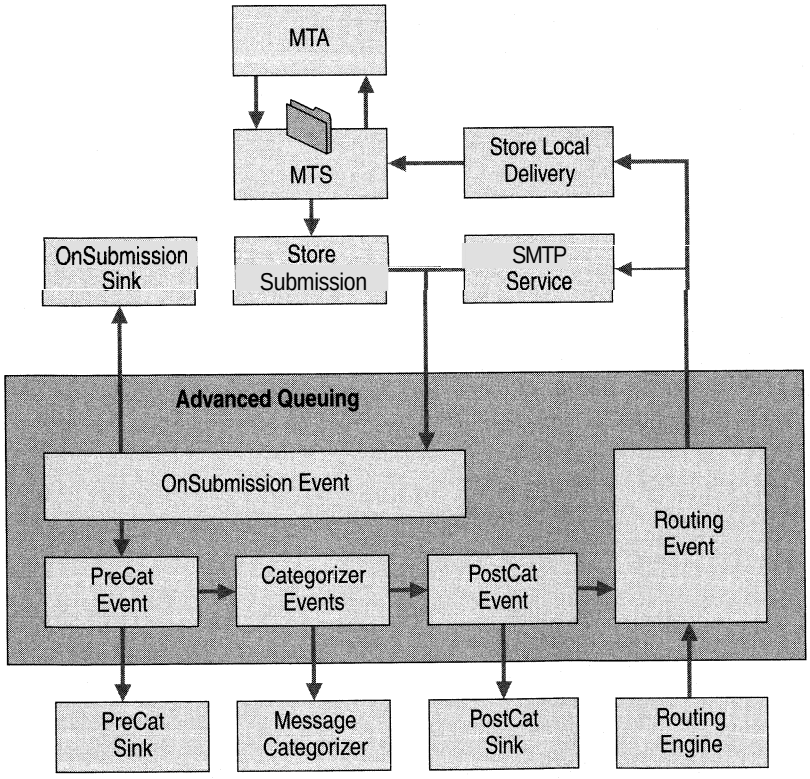


Figure 27.3 The event sink pathway in Web Storage System

Workflow Components For Exchange 2000

Workflow components provide you with greater control over the impact of application development on your servers. In the past, a developer could accidentally delete files, disrupt operations, and possibly bring the server down with the wrong code. Moreover, developers are often given broad access to objects on the server, which can raise security concerns. These concerns can be resolved in Workflow Designer.

Restricted Mode limits a developer's ability to compile objects. In this mode, no compiled objects are permitted. The developer's code has limited access for binding to or deleting objects on the server. This produces a much safer production environment. Privileged Mode permits developers to perform almost any function on your server, including using compiled objects.

Role membership determines who is eligible for a given mode. You grant this membership using the Component Services snap-in to manage Role membership on the Workflow Event Sink COM+ application. For more information about Restricted Mode and Privileged Mode, go to http://www.microsoft.com/exchange.

Workflow

Workflow describes applications that are modeled as business processes. Typical workflow applications include forms routing and approval, document review and publishing, and issue tracking. While you can implement such applications in nearly any programming language or development environment, you can simplify the task with the use of a workflow engine and specialized workflow modeling tools. Such tools allow the overall design, or flow, of a business process to be specified in a simple, high-level representation called a *process definition*. You can easily modify or extend the process definition without rewriting all of the low-level application code. The workflow engine executes and manages individual instances of a process definition, also known as process instances.

Creating Workflow Processes

Workflow Designer for Exchange simplifies application development. When you create a workflow process, the workflow designer creates a set of rules for the selected folder and its association with an event sink. The set of rules is called a process definition, the association is called an event registration, and the event sink is the CDO Workflow Event Sink.

The Workflow engine and event sink are provided with Exchange 2000 server. Workflow applications built for Microsoft Exchange 2000 Server use the object model to create process definitions consisting of a table of actions, or rules, that define the business process. At run-time, the applications call the workflow engine to create and manage process instances. The workflow engine evaluates the process definition to see which actions should be run, and in turn executes either VBScript code or compiled COM objects to perform application-specific business logic.

Using CDO Workflow Event Sinks

Exchange 2000 includes a Workflow Event Sink to write low-level code for security operations, the scripting environment, resource sharing, and data binding. The CDO Workflow Event Sink is the interface between Web Storage System process and the workflow engine. It is registered automatically as a COM application package when you install Exchange. You register the event sink in a folder, where it intercepts all changes made in that folder and automatically calls the workflow engine on your application's behalf to create and manage process instances. The event sink also hooks into system timer events to automate expiring items, overdue work, or other cleanup tasks. The event sink works well in scenarios where users are creating or editing documents using prepared applications such as Microsoft Word or Microsoft Excel.

Workflow Designer applications build on Web Storage System architecture, and must run on an Exchange 2000-based server. You typically install them into an application folder—either a public folder or private Inbox folder. The association between your folder and the event sink makes it possible for Web Storage System to notify the event sink when an event occurs in your folder. The event sink in turn calls the workflow engine to process the work item's transition.

Authorized folder owners register the Workflow Event Sink in application folders that they own. By granting the developer Restricted Mode or Privileged Mode, you control which users are authorized to register workflows on a per-server basis and which users can write "unlimited" workflow scripts.

> **Note** Microsoft Exchange Server 5.5 includes Routing Objects libraries, which are helper objects for building simple document routing applications. While similar to the Workflow Designer in some ways, the Routing Objects are designed only for MAPI-based e-mail routing applications. Exchange 2000 Server is compatible with existing MAPI applications, including Routing Objects. If you want to build a routing application that runs on both Microsoft Exchange 5.5 and Exchange 2000 servers, you should use MAPI and the Routing Objects libraries. If you want to write Web-based workflow applications or use Web Storage System features, you should use Exchange 2000 Server and Workflow Designer.

The ability to create workflow processes is an essential service provided by Exchange Workflow Designer. The workflow process is a series of tasks or actions, the order in which they must be performed, permissions defining which users can perform them, and script that is executed for each action. Exchange Workflow Designer provides a graphical view of the workflow process that you can easily create and edit from within the tool.

In its simplest form, the workflow process automates and enforces the order of the tasks in the process. For example, a user creates a new issue and assigns it to another user, who then resolves the issue and assigns it back to the original user, who then can close the issue. By defining the order of the tasks in the workflow process, you can ensure an issue is resolved before it is closed.

The conceptual model for a workflow process includes states and actions. A *state* is a discrete value of the state property of an item. For example, the state of an issue defines where that issue is in the workflow process, such as Resolved or Closed. An *action* defines the operations that can be performed on an item. One special type of action is a transition, which moves the item from one state to another. Actions are triggered by database changes.

Figure 27.4 shows a diagram of a workflow process used in the issue tracking example. The boxes represent states and arrows represent actions.

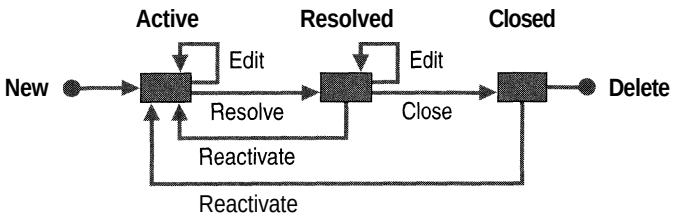


Figure 27.4 Workflow process for issue tracking example

Using script with workflow actions can make your workflow process more efficient. Before adding script to a workflow, it is important to understand the following:

- Workflow Engine
- Adding the Workflow System Account

Workflow Engine

The workflow engine controls the changes in a document's status in your workflow folder. It runs in response to certain system and Web Storage System events and encapsulates the logic for advancing the state of your workflow documents. Thus you can control the status of documents you are tracking or guiding.

Exchange 2000 includes an in-process server (CDOWF.dll) that implements the **IProcessInstance Advance** function of the workflow engine. The workflow event sink (CDOWFEVT.DLL) calls the engine when an event fires in your workflow-enabled folder. Figure 27.5 shows the workflow engine process.

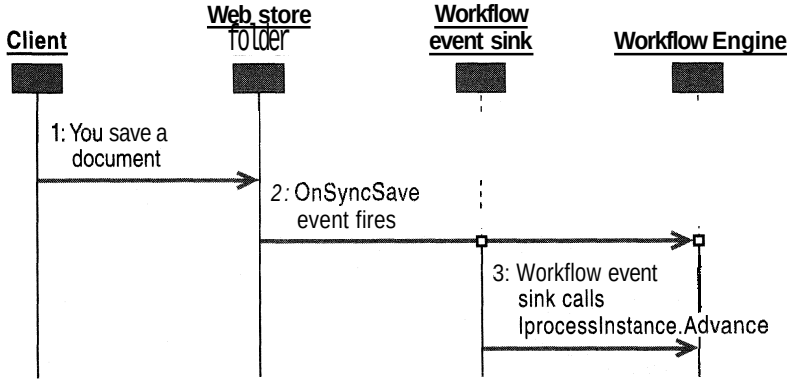


Figure 27.5 Workflow engine process

You do not call the engine directly unless you are writing an event sink. Other applications can use the workflow engine by making direct calls to the **IProcessInstance Advance** function. The event sink passes the engine one of the event types used in your process design as well as useful context information. The engine maintains document integrity based on the parameters the event sink passes.

You define a workflow process and set it up to be executed based on Exchange Web Storage System events or system events. The engine checks your process design when it receives an event message. The process design encapsulates the logic for advancing the state of your workflow documents. It includes script or COM actions and conditions.

The engine attempts to find a match between the message received and a rule in your process logic. When it finds a match, it executes the prescribed action. The engine and the script host run in the workflow event sink process and have access to the IProcessInstance and the user context.

When the workflow engine finds a row match and executes the row's action script or COM object, it passes the script host environment a **WorkflowSession** object with an ADODB Fields collection representing the ProcessInstance (row) undergoing transition. Only the workflow engine can change the status of a ProcessInstance directly. Figure 27.6 presents the workflow operation concepts.

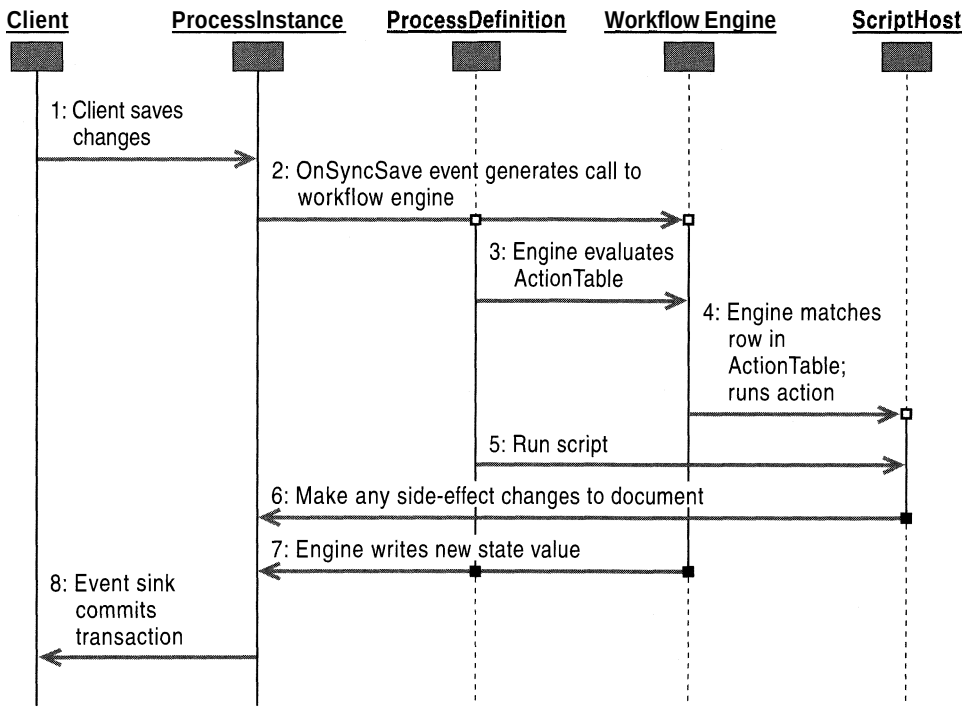


Figure 27.6 Workflow operation concepts

Adding the Workflow System Account

You need to add a Workflow System Account as a member of Domain EXServers group for the workflow event sink to run with proper privileges. You use Active Directory Users and Computers and Component Services to do this. Exchange 2000 Server Setup installs Workflow components by default.

To add the Workflow System Account

1. Pick an existing account from your Microsoft Windows 2000 domain or create a new account. These instructions assume that you have an account named Workflow System Account.
2. Set the account's password to never expire.
3. Make sure that the Workflow System Account has an Exchange mailbox. If your application uses e-mail-based workflow—that is, it uses the IWorkflowMessage::SendWorkflowMessage method—the Workflow System Account mailbox folder must be on the same physical server as the application folder where you are registering the workflows.

> **Note** The restriction relating to physical server location does not apply to the IWorkflowMessage::Send method, which uses SMTP.

4. Make the Workflow System Account a member of the Windows 2000 group named Domain EXServers. This enables the CDO Workflow event sink to execute with proper privileges.
5. Select your computer in the console tree of Component Services. Then select the COM+ application named **Workflow Event Sink.** Select **Properties** from the **Action** menu.
6. Select the **Identity** tab, click **This user,** and type the name and password for the Workflow System Account.
7. Use the Domain Controller Security Policy management console to configure the Workflow Event Sink account to act as part of the operating system. If you are configuring a member server (rather than a domain controller), use the Local Security Policy management console.
8. Add the Workflow System Account to the Privileged Workflow Authors role in the COM+ application package for the workflow event sink.

Workflow Security

When an action table includes scripts and COM objects that run on your server, security is an important issue. You need to prevent these scripts and COM objects from causing problems on your server and still allow users the flexibility to design their own workflows. The Workflow system account must have permission to act as part of the operating system, which is administered through Microsoft Management Console (MMC). Workflow Components solves this problem by implementing workflow security modes. Table 27.5 summarizes the two modes of security that Workflow Designer implements.

Table 27.5 Workflow Designer security modes

Factors to Consider	Restricted	Privileged
Script Access	Only the current Web Storage System item (ProcessInstance) is accessible	Scripts and objects can access enterprise databases as security context allows.
Security Context	EUSER_EXSTOREEVENT (guest privileges)	Workflow System Account defined by system administrator
CoCreateable COM objects	None	Unlimited registered components Allows you to integrate with other systems such as SQL databases and other business applications that provide COM components. Can use LDAP and Active Directory.

The security mode is a property of the process definition that you set at design time. You have the option of either setting this property to restricted or privileged. You use the restricted mode to limit workflows that are not trusted. The user needs to be able to set up simple document tracking applications, but you don't want them to run any script or COM object.

You must be a member of PrivilegedWorkflowAuthors group for your workflow to run in privileged mode. This is checked at run-time. PrivilegedWorkflowAuthors is a built-in group, and only you as the network administrator can add a developer to this group.

The workflow engine enforces security by checking privileges at run-time. For privileged mode workflows, it verifies that all the design-time pieces of a workflow are modified last by a member of the PrivilegedWorkflowAuthors role. This includes:

- ProcessDefinition
- CommonScript
- Event Sink Registration Document

The call to Advance fails if any of these are modified last by any users other than PrivilegedWorkflowAuthors.

Exchange 2000 Server setup creates the COM+ Application package for the Workflow Event Sink and installs the PrivilegedWorkflowAuthors role. If you write your own event sink, you must register it as COM+ Application package in the Component Services tool, and add the PrivilegedWorkflowAuthors role. Your sink's calls to the workflow engine will fail if this role is not present in the sink's security context.

Application Development Tools

Following are some of the Exchange 2000 development tools that you can use to develop applications.

Web Storage System Viewer

The Web Storage System Viewer allows you to see the contents of folders in Web Storage System.

For more information about this tool, see "Messaging and Collaboration" at http://www.microsoft.com/exchange.

Web Storage System Schema Designer

The Web Storage System Schema Designer is a tool that allows you to view and modify the schema for folders in Web Storage System. This tool is written in Visual Basic and you can run it on both the client and the server. It uses WebDAV to communicate to the server to retrieve and update the schema, While users of the tool should have an understanding of how the schema is supported in Exchange, they do not need to know the details of how schema information is stored.

For more information about Web Storage System Schema Designer, go to http://www.microsoft.com/exchange.

The Web Storage System Application Deployment Wizard

The Web Storage System Application Deployment Wizard simplifies deploying applications. The application packaging scans a virtual root and builds a script that allows you to deploy the application on another server or group of servers. The distribution plays the script and builds the application on another server.

For more information about Web Storage System Application Deployment Wizard, go to http://www.microsoft.com/exchange.

The Event Sink Template Wizard for Visual Basic 6.0

The Event Sink Template Wizard for Visual Basic 6.0 creates a Visual Basic project that implements the necessary components needed to create the DLL. This assists developers in building event sink DLLs in Visual Basic.

For more information about Event Sink Template Wizard for Visual Basic 6.0, go to http://www.microsoft.com/exchange.

CHAPTER 28

Backup and Restore

Your organization can be prepared to recover from data loss by developing a consistent backup and restore strategy. Implementing that strategy and maintaining database consistency can improve the integrity of any Microsoft Exchange 2000 Server database. It is important to understand the Exchange 2000 database technology and the different categories of backup that are available. Knowing the prerequisites and processes involved will help you to back up data to prevent data loss and to restore data when needed.

In This Chapter

Exchange 2000 Database Technology

Backup Prerequisites

Backup Categories

Backup Process

Restore Process Overview

Prerequisites for Disaster Recovery

Restore Categories

Restore Constraints

Database Consistency

Recovering from Disasters

Best Practices

Exchange 2000 Database Technology

Important Before you back up or restore your Exchange 2000 data, you should download the latest hot fix from the Microsoft Exchange 2000 Server Web site at http://www.microsoft.corn/exchange to get the most recent version of Windows 2000 Backup (NTBackup.exe).

Understanding the underlying database technology that Exchange 2000 uses can help you to better understand the backup and restore process. The database technology is based on the Extensible Storage Engine (ESE), which is part of the Microsoft Web Storage System process. Users of the affected database will not be able to access their e-mail during a restore. The ESE also has features that allow Exchange 2000 to service user accounts while the server is being backed up. Figure 28.1 shows a current database.

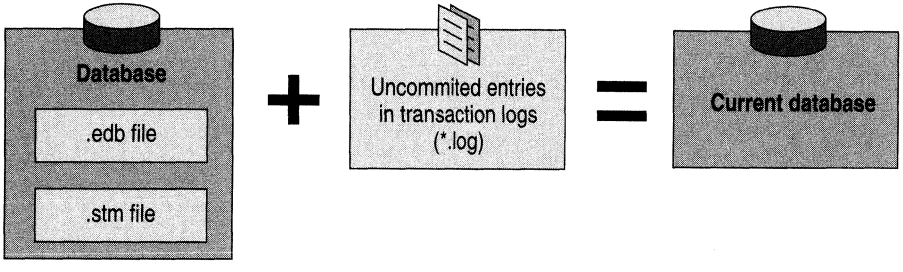


Figure 28.1 Description of a current database

Exchange 2000 uses Microsoft Windows 2000 Active Directory. In earlier versions of Exchange, the directory was an ESE database that was backed up with other Exchange databases. Because the Exchange 2000 directory is part of Windows, backing up Windows 2000 Active Directory is as important as backing up Exchange 2000 databases.

Extensible Storage Engine Structure

The ESE uses a balanced-tree (B-tree) structure to store data. Each page in the database file is a node in the B-tree structure. An ESE database can contain up to 2^32 pages or 16 terabytes (Active Directory uses 8-kilobyte [KB] pages and can contain up to 32 terabytes). This means that your database size is limited only by your ability to back up and restore the database in a timely way.

Multiple Storage Groups and Failover

An Exchange storage group is a part of the ESE. Each storage group consists of the properties database (.edb) files, streaming database (.stm) files, and the set of log files for all the databases in the storage group. Exchange 2000 provides multiple storage groups and multiple databases within each storage group. Failover is a cluster service event where the active node stops functioning, or a critical service on the active node fails, and the secondary node takes its place. When you install Exchange 2000 on each node, you use Cluster Administrator to define the failover rules for Exchange 2000 services. These rules define the conditions that trigger a failover to the secondary node. The node that takes over must have access to the data at the moment of failure, so users and jobs have current data.

Databases

You can create multiple databases within each cluster service storage group and restore databases in a storage group individually or together. A single set of log files contains entries for all databases in the storage group. One database within the storage group is composed of the following files:

- The .edb file, which contains folders, tables, and indexes for messaging data, and MAPI messages and attachments.
- The .stm file, which is a new format in Exchange 2000 for storing native Internet content. The internal schema for the .stm pages is stored in the .edb file. The .edb file stores messages in Rich Text Format (RTF). The .edb and .stm files function as a pair, and the database signature (a 32-bit random number combined with the time the database was created) is stored as a header in both files.
- Site Replication Service (.srs) files, which permit compatibility with Exchange 5.5 by emulating an Exchange 5.5 directory service.
- Key Management server (.kms) files, which provide security encryption services.

The .edb and .stm files are organized into 4-KB pages. In the .edb file, the pages are organized into a B-tree. However, in the .stm file, the pages are grouped in a clustered block or run format, similar to a file system such as NTFS. These pages are stored in 64-KB blocks (16 runs of 4 KB each) within the .stm file on the hard disk.

Exchange data is stored in .edb, .stm, .srs, and .kms files. These are the database files that share the same attributes as the operating system files. The .edb files are organized into a collection of pages that store data definitions, data, and indexes. These pages are numbered sequentially in a database file.

Logs form in the following way:

1. When the current log file fills, the Edbtemp.log file starts.
2. The original log file name changes to the true generation number, so it has the Edb prefix followed by a five-digit hexadecimal number. If the first generation file name is Edb000A1.log, the next Edb.log file name is Edb000A2.log.
3. The file Edbtemp.log becomes Edb.log and the process repeats.

Note Exchange 2000 uses the base name, not the Edb file name prefix. They are presented in this format for purposes of clarity. The actual file base names appear differently in your directory.

Database GUID

A database globally unique identifier (GUID) is assigned to the store database and a matching GUID is stored in Active Directory. This is important, because the database does not mount if the GUIDs do not match.

Mailbox GUID

Each mailbox in the database has a GUID. The user account in Active Directory has the GUID of the mailbox that it owns, which means that even though you can delete a user from Active Directory, the mailbox still exists in the database. Therefore, you can reconnect the same mailbox to a different user. If the user account is accidentally deleted, you can fix it by reconnecting the mailbox. It adds the mailbox GUID as a property field to the user account in Active Directory. The mailbox itself has a default retention time of 30 days. For more information about how to recover a single mailbox, see "Restoring a Single Mailbox" later in this chapter.

Transaction Log Files

Changes to the database are also stored in transaction log files. These files store a sequential list of every operation that is performed on a page in memory. These log files give you a way to restore committed transactions that you can lose when there is a power failure or some other type of disaster. Exchange 2000 log files are always 5 megabytes (MB) in size. If you view them in Windows Explorer, they show up as 5,242,880 bytes. If the log files do not show this exact size, typically they are damaged.

Each storage group service that uses the ESE reserves two log files, Res1.log and Res2.log, which are stored in the log directory. Reserved log files are used as placeholders for extra disk space that can be used if the service runs out of space. They are the same size as, and function like, normal log files. It is recommended that you store log files and database files on separate drives.

Log File Signature

You can use the log file signature to your advantage. If you shut down the server and inadvertently delete all log files, when you restart the server, the ESE creates a new series of log files, starting with a generation number of one. Because log files can have the same name, the ESE stamps the header in each file series with a unique signature so it can distinguish between different series of log files.

Checkpoint File

Checkpoint files store information that indicates when a transaction is successfully saved to the hard disk. The Edb.chk file points to the log file that has had all transactions successfully committed to the database file. When all the pages that are changed by transactions in a given log file write to the database file, the checkpoint advances to the log file with the next unwritten entry. Separate Edb.chk files are maintained for each storage group. The Edb.chk file is not required if you want to replay transactions. The ESE determines which transactions are already written by examining the transaction log files. However, using the checkpoint saves the ESE from starting at the first log file and checking every operation.

Figure 28.2 shows a description of checkpoint files.

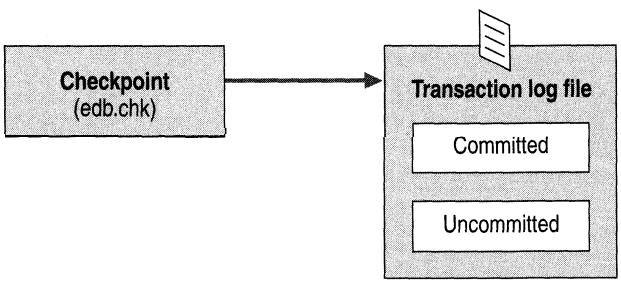


Figure 28.2 Description of checkpoint files

Circular Logging

Circular logging reduces disk storage requirements by overwriting transaction logs after transactions are committed to the databases. Circular logging is used with non-critical data and is disabled by default. Transactions to the .stm files are not logged because the data is saved to the .stm file before the transaction is committed. Because less data is written, the process is faster and more resources are available. When circular logging runs, changes to the .edb file are written to the transaction log files. In a restore operation, you can play back data only to the point of the last backup. These transactions are entered into the restored database file to update the database to the point in time when the last backup was performed, and not to the point when the database files went offline. It is recommended that you leave circular logging disabled because you can lose information during a restore operation.

Checksum

The checksum (also called a message hash) is a string of 4-byte bits that is calculated and then added to the page. Every .edb file is made up of 4-KB pages and the integrity of these pages is verified through a checksum and a 4-byte page number in the header of the database page. The first 82 bytes of the database page contain the header information, which contains flags for the type of page and information about what kind of data the page contains. When Exchange 2000 reads pages out of the database, the pages are compared for the correct page number and for the checksum, which is calculated to ensure that the data in the page is undamaged. If the data is damaged, the ESE returns an error to Web Storage System, the database is stopped, and an event is logged informing you of damage and actions to take (for example, completing a restore). This process ensures the optimal integrity of the database by informing you immediately if your database is damaged.

Indiv Database Backup

In Exchange 2000, you can use the backup tool included with Windows 2000 to perform an individual backup. It is recommended that you download the latest hot fix or service pack, which includes this tool, from the Microsoft Exchange 2000 Server Web site at http://www.microsoft.com/exchange. Backup recognizes multiple storage groups and multiple databases. The level of detail for backup and restore is determined at the database level. It is possible to restore an individual database inside a storage group without restoring an entire storage group. If you tag an individual database for backup, it backs up the .edb and .stm files and the needed transaction log files. As part of your backup strategy, it is recommended that you tag an entire storage group for backup, because backing up databases individually also backs up log files multiple times.

Exchange 2000 Server Clusters

Exchange 2000 supports active/active clustering. A cluster is a group of two or more independent servers (nodes) that are managed as a single system. These nodes have individual memory, processors, network interface cards, and a shared storage medium. The nodes must also have identical processors and the same amount of RAM. You connect these computers using a network cable that puts them in continuous contact with each other. If one node fails, the other nodes of the cluster take over the failed node's clients. There may be a short period of time during which clients cannot connect to the network, but when the hand-off is completed, normal working conditions resume.

One or more Exchange virtual servers exist in the cluster and each virtual server runs on one of the nodes in the cluster. Exchange can support multiple virtual servers on a single node. Clients connect to the virtual servers the same way they would connect to a stand-alone server.

A single Exchange 2000 server can support up to four storage groups, with five databases per storage group. Each storage group has a temporary database containing temporary table5 and is used for large sorts. Exchange also holds 12 storage groups in reserve for restore. For a server cluster, Web Storage System can mount four storage groups. If one of the nodes stopc functioning, the other node must be able to support the storage groups of both nodes. Figure 28.3 shows the files found in one storage group.

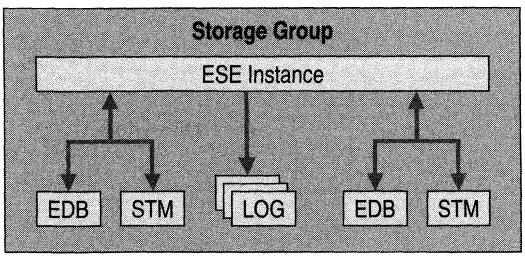


Figure 28.3 Storage group files

The unit of failover is the Exchange cluster group (also called the Exchange virtual serverj. The Cluster service, not Exchange, takes over in a failover event. Clustering resources, such as MSExchangeIS and MSExchangeMTA, are grouped together into a cluster group through Cluster Administrator. When one of these resources fails, the entire group fails because these resources depend on each other. Then the Cluster service moves the group to another node in the cluster. There might be a brief interruption in service during a failover event, although clients such as Outlook usually reconnect automatically.

Exchange clustering provides high availability for messaging. The Cluster service monitors the virtual servers in the cluster. In the event of a failure, the Cluster service restarts or moves the affected virtual servers to a healthy node. For planned outages, you can manually move the virtual servers to other nodes. In either event, the client experiences an interruption of service only during the brief time that the virtual server is offline. The ability to use multiple servers at all times reduces system costs while increasing reliability, because no dedicated failover-only or hot standby servers are needed.

Clustering support differs according to the component. Table 28.1 details which components are supported, and in some cases, the type of clustering they can support.

In Table 28.1 the following terms are used to describe the level of functionality that each component supports.

- *Active/passive*. Only one instance of the component can run in a cluster.
- *Active/active*. Multiple instances of the component can run in a cluster.

Table 28.1 Clustering support

Components	Notes	Cluster Functionality
Exchange System Attendant		Activelactive
Web Storage System	After failover: each node can support four storage groups (with more in reserve).	Active/active
Message Transfer Agent (MTA)	One instance per cluster. The MTA is in only one Exchange virtual server.	Active/passive
POP3 (Post Office Protocol version 3)	Multiple virtual servers per node.	Active/active
IMAP Internet Message Access Protocol)	Multiple virtual servers per node.	Active/active
SMTP (Simple Mail Transfer Protocol)	Multiple virtual servers per node.	Active/active
Chat		Active/passive
MSSearch		Active/active

Backup Prerequisites

To successfully back up the data on your Exchange 2000-based servers, you must satisfy the necessary prerequisites. You need to define a backup strategy, fulfill the resource requirements, have the proper Web Storage System, define the data you need to back up, verify the integrity of your configuration data, and address performance considerations.

Planning a Backup Strategy

Your backup strategy determines your restore strategy. These operations cannot be planned separately. When you create a backup strategy, you should also consider the time it takes to restore. Make sure you have enough hard disk space to restore both the database and the log files. For example, a weekly full backup plus one week of transactional log files might be more than your server can store, depending on how many log files are generated during a week. If your server generates 2,000 log files in a week, that amounts to 10 gigabytes (GB) of log file space, in addition to the database.

Creating and Verifying Daily Backups

Creating and verifying daily backups is a critical step for successful restores. You can verify with practice restores and then examine the detailed log files using Windows 2000 Backup (NTBackup.exe). However, you can only restore valid data if the backup is valid. Failure to verify backups is a common mistake because it is easy to assume that backup tapes are swapped and that data is backed up properly. Make it part of your daily routine to review all backup logs and to follow up on any errors or inconsistencies.

Standardizing Backup Formats

Restore equipment must be compatible with production equipment. If you deploy a new type of drive, make sure that you use a compatible model for your restore equipment. You should also test reading and restoring production backups on the drive you use for restore.

Publishing a Maintenance Schedule

It is important to set user expectations by publishing a maintenance schedule, especially when users expect service 24 hours a day, seven days a week. Maintenance is inevitable because of the nature of data processing upgrades.

Checking for Hot Fixes

Microsoft releases hot fixes to correct problems that might occur, so check the Microsoft Web site at http://www.microsoft.com monthly for the latest hot fixes.

Deciding the Types of Data to Back Up

It is important to back up mission-critical data, which is typically the user's data. This might include configuration information, mailboxes, or databases, depending on your organization's priorities. You should always have a full backup of your static data and routinely back up your dynamic data.

Static Data

Static data includes the following types of information:

- Windows 2000 operating system and any service packs or hot fixes
- Packaged application software (for example, Exchange)
- Supporting software, such as third-party backup software or system management software
- User application software, such as Active Server Pages (ASP) applications, mailbox agents. and workflow software
- Management scripts

Dynamic Data

Dynamic data includes the following types of data:

- Exchange Web Storage System databases and supporting files
- Message tracking log
- Active Directory
- Key Management Service databases (need to run on a member server with certificate authority)
- Site Replication Service (SRS) databases
- System state including Internet Information Services (IIS) metabase (and quorum if on a cluster)

> **Note** The composition of system state varies, depending on whether you use Windows 2000 Professional or Windows 2000 Server. Both versions of Windows 2000 have system states that include: boot files, COM+ Class Registration database, and the Registry. But Windows 2000 Server System State also includes: Active Directory (if the server is a domain controller), Sysvol (if the server is a domain controller), Certificate Services database (if this is run), and quorum (if the server is part of a cluster).

Resource Requirements

Before you back up your server, you must ensure that you have adequate resources. Table 28.2 can help you calculate your annual backup requirements. By filling in the boxes and summing the results, you can create a Backup Inventory plan. Table 28.2 helps you consider all aspects of tape use: accident, backup, disaster restore, and replacement.

Table 28.2 Backup tape inventory

Number of backup drives		Number of tapes to back up a set		Number of sets in a rotation schedule		Annual number of rotation schedules		Annual number of tapes required
	X		X		X		=	

Add the number of tapes you think you will need to replace and the number of tapes you expect to retire to your annual total.

Exchange 2000 Services and Permissions

To back up Exchange 2000, you must have Active Directory available and the Web Storage System services running on the server. If one of the services is not running or is unresponsive, the backup fails.

To back up data, you must have Backup Operator permissions on the local computer. Windows 2000 Backup uses the permissions of the current logon to do the backup. Third-party backup utilities can function like Windows 2000 services, which use permissions from the service startup parameters. These are typically the permissions set in the LocalSystem account.

Active Directory

While Active Directory is not a component of Exchange 2000, it is relied upon heavily. Configuration information is stored in Active Directory. User objects are kept in Active Directory and are mail-enabled for Exchange.

Internet Information Services

While IIS is not a component of Exchange 2000, it is relied upon heavily. Some Exchange configuration information, such as Internet Protocol (IP) configuration and message routing information, is stored on the local computer. For a successful restore, you must back up the IIS metabase. The metabase is a structure for storing IIS configuration settings. You can view the metabase using utilities such as MetaEdit and Mdutil.

The metabase update service is a component in Exchange 2000 that reads data from Active Directory and transposes it into the local IIS metabase. It runs underneath the Microsoft System Attendant Service and accepts Active Directory notifications about changes in the directory that pertain to the configuration of protocols that reside in IIS. When it receives these notifications, it updates the metabase automatically.

This service allows you to make remote configuration changes to virtual servers without a permanent connection to each system.

With Windows 2000 Backup you can back up files on the local computer or on network drives, in addition to backing up the system state or exchange databases (if Esebcli2.dll is installed). A file system backup refers to backing up files. You can select an entire hard drive and then back up the entire directory, or you can select an individual file and then back it up on a per file basis. If the file is locked, the backup skips the file.

File system backup is not the best choice for the metabase because the metabase maintains dependencies on other components that are not saved using a straight file system backup. Also, the backup file might be undergoing modifications at the time of backup. The best method of backup is to perform a system state backup, because this backs up the metabase. The Metabase.bin file is skipped and not retrieved by a file system backup if IISAdmin is running.

If you only retrieve the Metabase.bin file in a file system backup, that is not enough to restore it if you need to rebuild the server. The Metabase.bin file is encrypted and for IISAdmin to open it, the public keys for the system also need to be backed up and restored. Public key information is in both the registry and in the Documents and Settings\Administrator\Application Data\Microsoft\Crypto folder. It is recommended that you back up the metabase using a system state backup because of these dependencies.

Performance Considerations

You should complete backups in a timely way, and certainly before the next backup begins. For most organizations, performance is measured by the length of time for the completion of the restore process. You should consider using the fastest hardware possible. Connecting the tape drive to the server itself further speeds up the restore because the network does not slow it down. Consider a tape drive that uses streaming and striping. Alternatively, consider using a fast, wide SCSI connection. The more tape drives involved, the faster the speed in writing to backup.

If the hard drive on which Windows 2000 is installed is mirrored, the transaction logs are placed on a dedicated physical hard drive, which can also be mirrored. If the Windows 2000 swap file and Web Storage System are placed on a mirrored drive, backup can require up to twice the disk space of the database file. Therefore you should make sure to allow sufficient disk space.

Capturing Configuration Data

To capture configuration data, it is best to perform a full backup periodically. A full backup should include the system state and all Windows and Exchange binary files. You might want to perform this backup during scheduled maintenance. A full offline backup is not required for backing up Web Storage System. It is recommended that you use online backups to back up Web Storage System.

Performing Active Directory Backups

It is recommended that you also back up Active Directory information on your domain controllers. Performing a system state backup on your domain controllers captures Active Directory and the Sysvol directory. You should back up Active Directory and Exchange 2000 databases at the same time to avoid losing configuration or user objects on the domain controllers or global catalogs. It is important to back up your domain controllers because it allows you to quickly restore all your domain controllers in parallel instead of having to wait for replication to update them.

Backup Categories

The type of backup you need to perform varies depending on the importance of the data you store. Each of the backups categories has advantages and disadvantages in terms of data storage, performance, and time requirements. The backups discussed are:

- Full backup
- Copy backup
- Incremental backup
- Differential backup
- Backing up data on a cluster server node

For more information about backup categories, see the *Microsoft Windows* 2000 *Sewer Resource Kit.*

Full Backup

A full (or normal) backup backs up the entire Web Storage System and the Exchange log files. It deletes the transaction log files that contain transactions that are committed to the server database. This means that daily full backups prevent log files from consuming space on the hard disk. Restoring from a full backup requires only the full backup file. This is generally the preferred backup method.

Copy Backup

A copy backup is the same as a full backup except that it does not delete log files. You can create a copy backup if you want a full backup without disturbing the state of other ongoing full, incremental, or differential backups. If you plan to install new software or implement a system change, you should create a copy backup.

Incremental Backup

An incremental backup backs up log files prior to the checkpoint log file, and then deletes them. An incremental backup also backs up all transaction log files and deletes the log files that contain transactions that are committed to the database. This type of backup cannot be used when circular logging is enabled. Restoring from an incremental backup requires that you have the last full backup and each subsequent incremental backup. After this is complete, the transaction logs are applied in sequential order to the Exchange database that was restored with the full backup. It is important that all incremental backups are restored prior to starting log file replay. Otherwise, you might lose data.

Each backup is the same size, given the same database activity. To restore the database, you need the last full backup and all subsequent incremental backups. This is the fastest backup process.

If any incremental backup is damaged (that is, if any log file in the set is damaged), you cannot restore incremental backups made later. Remember that transactions must be replayed in order. One damaged log file prevents subsequent log files from playing again. The ESE enforces this to prevent damage to the database.

Differential Backup

A differential backup backs up log files prior to the checkpoint log file, but does not delete them. Each backup becomes increasingly larger in size. A differential backup backs up all transaction log files, but does not delete them. Restoring from a differential backup requires that you have the last full backup and the most recent differential backup. This is the next fastest restore process after a full backup.

Backing Up Data on a Cluster Server Node

You should use Exchange 2000 Backup Wizard to perform a backup of a cluster node where the Cluster service is operational. In the **What to Back Up** dialog box of the Exchange 2000 Backup Wizard, click **Back up everything on my computer.**

- Be sure the node on which you perform the backup is the owner of the cluster quorum disk. To check this, stop the Cluster service on all other nodes except the node running Windows 2000 Backup (NTBackup.exe). Then choose between the following options:
  - Click **Back up everything on my computer** to back up everything on a node. This backup includes the clustering software, the cluster administrative software, the quorum, and the system state data.
  - Click **Only back up the System State data** to back up the system state, which includes the quorum.
- To back up all cluster disks owned by a node, perform the back up from that node.

  Note During backup, Windows 2000 Backup might report an error that says, "Completed with Skipped Files, Examining the Windows 2000 Backup log, both CLUSDB and CLUSDB.LOG failed to be backed up." Ignore this error. The quorum logs from the cluster quorum drive are successfully backed up.

After backing up the cluster quorum disk on one node, you don't need to back up the quorum on the remaining cluster nodes. You can also back up the clustering software, the cluster administrative software, the system state data, and other cluster disks on the remaining cluster nodes.

Backup Process

You begin the backup process by starting the backup application. You make calls to Web Storage System with the type of backup desired and then the backup procedure begins. Web Storage System informs the ESE that it is entering a backup mode, and then a patch file is generated for each database in the backup (assuming this is a full backup). If this is a differential or incremental backup, it does not create a patch file. When the ESE enters a backup mode, a new log file opens. For example, if Edb.log is the current open log file, Edb.log is closed and is renamed to the latest generation and a new Edb.log is opened. This indicates the point when the ESE can truncate the logs, after the backup is complete.

When the backup begins, the agent requests that the database read and sequence the database pages from the ESE. The pages are read in numeric sequence in groups of 16 4-KB pages (though the actual size can vary). As the database engine reads the pages, the ESE verifies them through a checksum to ensure that they are valid. If they are invalid, the backup stops to prevent the storage

of damaged data. After the backup is complete and all the pages are read, the backup copies the logs and patch files to the backup set. The log files are then truncated or deleted at the point when the new generation started at the beginning of the backup. The backup set closes, the ESE enters normal mode, and the backup is complete.

In an incremental or differential backup, only the log files are affected. Operations that involve patch files, checksums, or reading pages sequentially are not executed. Figure 28.4 shows the Exchange 2000 backup process.

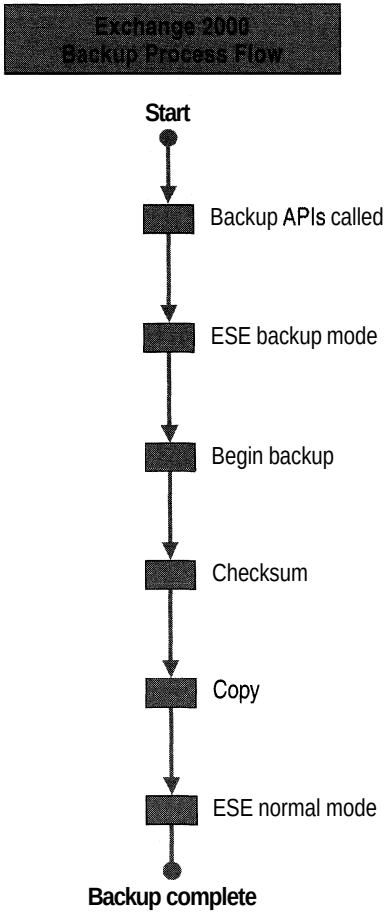


Figure 28.4 The Exchange 2000 backup process

Online Backup Process

There are two ways to perform backups: online and offline. It is recommended that you use online backup because it allows the database to continue running while you are backing up data. The advantage of an online backup is that users are not affected and no processing jobs are interrupted.

An online backup can be either partial or full. Full backups copy everything in the database and partial backups copy only the log files. Regardless of the category of backup performed, an online backup process always follows these steps:

1. The backup starts, a synchronization point is fixed, and an empty patch file is created.
2. Edb.log is renamed to the next log number regardless of whether it is full, and a new Edb.log is created.
3. Windows 2000 Backup copies 64 KB of data at a time. Additional transactions are created and saved as normal. Each page is checked for damage.
4. Logs used during the backup process (those from the checkpoint forward) and the patch files are copied to tape.
5. Backup finishes.

Table 28.3 Online backup types

Category	Data Backed Up	Logs Purged	Restore
Full (normal)	Database and log files	Yes	Last Normal
			Start LogFileReplay
Incremental	Log files only	Yes	Last Normal
			Every Incremental
			Start LogFileReplay
Differential	Log files only	No	Last Normal
			Last Differential
			Start LogFileReplay
Copy	Database and log files	No	Last Copy
			Last Differential or every Incremental
			Start LogFileReplay

During an online backup, transactions can still be entered in databases. If a transaction causes a split operation (an internal low-level ESE operation) across the backup boundary (the location in the .edb file that designates what has been and has not been backed up), the affected page before

the boundary is recorded in the patch (.pat) file. A separate .pat file is used for each database that is backed up, such as -PRIV1.pat, Pub1.pat, and Srs.pat. These files are seen only during the backup and restore processes. The steps are as follows:

1. The backup for the current storage group .edb file begins.
2. A .pat file is created for each database that is backed up and the database header is written into the .pat file.
3. Split operations across the backup boundary are written into the .pat file.
4. The .pat files are written to the backup tape with the log files.
5. The .pat files are deleted.

Offline Backup

Offline backups allow you to save a copy of a consistent database file. Offline backups are always full backups because the database shuts down. An offline backup is always the second choice, with the online backup being preferred. It is necessary to dismount the database before performing the offline backup.

Backup Rotation Schedule

Creating and verifying daily backups is a critical step in successful data restoration. Don't assume that backup tapes are being swapped and that data is being properly backed up. Make it a part of your daily routine to review all backup logs, and then investigate any errors or inconsistencies.

> **Note** The option to verify backups in Windows 2000 Backup is not a recommended practice. When Windows 2000 Backup writes data to the tape, it includes Cyclical Redundancy Checking (CRCs) within the data. After the backup is complete, Windows 2000 Backup reads the tape, verifying those checksums. The best way to verify your data is through a practice restore.

A good rotation schedule is important for data restore. The best schedule is one that provides you with a comprehensive history of file versions. Records should be stored in a secure, off-site location. Following are descriptions of two popular rotation schedules. For more information about backup schedules, see the *Windows* 2000 *Sewer Resource* ***Kit.***

Month-Week-Day

The most commonly used media rotation schedule is month-week-day. This scheme uses daily, weekly, and monthly backup sets. Backup media are labeled with the day of the week that it backs up. Typically, incremental backups are performed on the day group of media. This media is reused each week on the day matching its label. A set of up to five weekly backup media is labeled week1, week2, and so on. Full backups are recorded weekly, on the day that a day media is not used. This week media is reused monthly. The final set of the three media is labeled month1, month2, and so on, according to which month of the quarter in which they are used. This month media records full backups on the last business day of each month and is reused quarterly. Each of

these media can be a single tape or a set of tapes, depending on the amount of data you need to back up. A total of 12 media sets are required for this basic rotation scheme, allowing for a history of two to three months. Because a longer history is often required, archive tapes are periodically pulled from the rotation and replaced with new tapes. Table 28.4 shows the month-week-day rotation schedule.

Table 28.4 The month-week-day rotation schedule

Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Sunday
1 Day 1 incremental	2 Day 2 incremental	3 Day 3 incremental	4 Day 4 incremental	5 Day 5 incremental	6 Week 1 normal	7
8 Day 1 incremental	9 Day 2 incremental	10 Day 3 incremental	11 Day 4 incremental	12 Day 5 incremental	13 Week 2 normal	14
15 Day 1 incremental	16 Day 2 incremental	17 Day 3 incremental	18 Day 4 incremental	19 Day 5 incremental	20 Week 3 normal	21
22 Day 1 incremental	23 Day 2 incremental	24 Day 3 incremental	25 Day 4 incremental	26 Day 5 incremental	27 Week 4 normal	28
29 Day 1 incremental	30 Day 2 incremental	31 Day 3 incremental	1 Day 4 incremental	2 Day 5 incremental	Month 1 Week 1 normal	4

Towers of Hanoi

The Towers of Hanoi rotation scheme is widely used. It is based on a mathematical puzzle of the same name. This schedule follows the solution to this problem. One media set, A, is used every other backup session (daily). Start day 1 with A and repeat every other backup (every other day). Set B then starts on the first non-A backup day and repeats every fourth backup session. Set C then starts on the first non-A or non-B backup day and repeats every eighth session. Set D then starts on the first non-A, non-B, or non-C backup day and repeats every sixteenth session. Set E alternates with set D.

With each additional set added to the rotation scheme, the backup history doubles. The frequently used sets have the most recent copies of the data; while less frequently used sets retain older versions. You can use this schedule either on a daily or weekly rotation scheme. You should base the decision regarding the frequency of rotation on the amount of log file transactions. To maintain the required history of file versions, you should use at least five sets in the weekly rotation schedule, or eight for a daily rotation scheme. You should periodically remove tapes from the rotation for archive purposes. Table 28.5 shows the Towers of Hanoi rotation schedule.

Table 28.5 Towers of Hanoi rotation schedule

Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Sunday
1 A	2 B	3 A	4 C	5 A	6 D	7 A
8	9 A	10 B	11 A	12	13 A	14
15 A	16	17 A	18 B	19 A	20 C	21 A
22	23 A	24	25 A	26 B	27 A	28
29 A	30	31 A	1	2 A	3 B	4 A
5 C	6 A	7 D	8 A	9	10 A	11 A
12 A	13	14 A	15	16 A	17	18 A

Restore Process Overview

Before you begin the restore process, the database or storage group must be dismounted and made inaccessible to users. You can do this by using System Manager. When a restore operation begins, the store informs the database engine and the Extensible Storage Engine (ESE) enters restore mode. The agent (or the backup application programming interfaces [APIs]) copies the database (.mdb) from the tape directly to the database target path. The database is a file pair of the .edb and .stm files. The associated log and patch files are copied to the server in a temporary location specified by the backup operator so they aren't saved to the same location as current files in the Exchange or Production Database directory. If this happens, you can over-write log files and cause a logical corruption of the database. The restore instance of the database brings back the log and patch files to a temporary location specified by the backup operator and then restores those files to the server. After the files are restored, a new restore storage group (an instance of the ESE) starts specifically for the purpose of restoring the database. The patch file data and the log files from the

tape backup are then applied to the database by the restore database engine. The database engine processes the current logs, bringing you to the exact restore instance you want. After this is complete, some cleanup is done by deleting log and patch files from the temporary location and deleting the restore storage instance, and the storage group is mounted by the instance that owns it. Figure *28.5* shows the Exchange *2000* restore process.

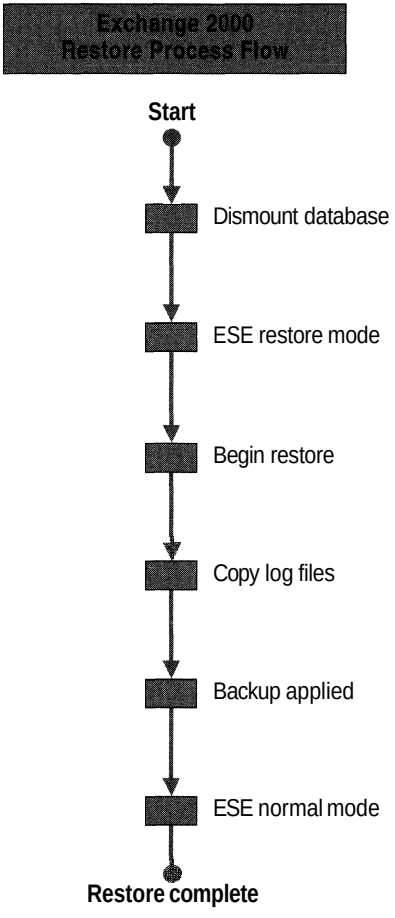


Figure 28.5 The Exchange 2000 restore process

Prerequisites for Disaster Recovery

Restore Windows 2000 using the CD and attach the computer to the network with a suitable tape drive. Ensure that the JP address and domain name are assigned to the new system before installing Exchange 2000. From that point on, you can restore data from backups. Early data restore efforts focus on restoring the operating system(s) for each computer system. Later, you can do first-line restores of application and user data from the backup tapes. Note that logs backed up from tape are not restored to the log directory. Instead, the user determines the location to which the logs are restored. The following are conditions you should note before restoring data:

- If the Exchange 2000 server to which you are restoring files is a member server in a domain, be sure that Active Directory is running. Run Exchange System Manager and verify that a valid server object still exists for the Exchange 2000 server in Active Directory. If Active Directory does not exist, restore Active Directory prior to restoring Exchange 2000.
- If the Exchange server you want to restore is also the domain controller, begin by restoring Active Directory on that computer. You can only restore Exchange 2000 after Active Directory is successfully restored. The security ID on the restored server must match the security ID of the original server. If the security IDs do not match, you cannot access Web Storage System until you restore only Web Storage System and then manually rebuild the Windows 2000 accounts.

Planning Restore Strategies

The backup strategy you choose has a direct impact on what you will do for restore. In a strategy with a full daily backup, only one tape is required. Thus, the volume of data and the time this procedure takes are minimized. However, in a full backup with daily incremental backups, you have the full backup tape from the first day of the cycle and a tape for each daily incremental backup until the point at which you wish to recover. With multiple tapes, there is a more data to manipulate and the restore takes longer. Unfortunately, the more tapes you have to install and the longer the restore takes, the more opportunities there are for you to make a mistake. The third strategy is a full daily backup with a daily differential backup. Two tapes are required: the full backup from the start of the cycle, and the last differential from the day before the point at which you want to recover. Most clients use a full daily backup strategy because of its convenience and speed. Figure 28.6 shows these three commonly used restore strategies. For more information about planning restore strategies, see the ***Windows*** 2000 ***Server Resource Kit.***

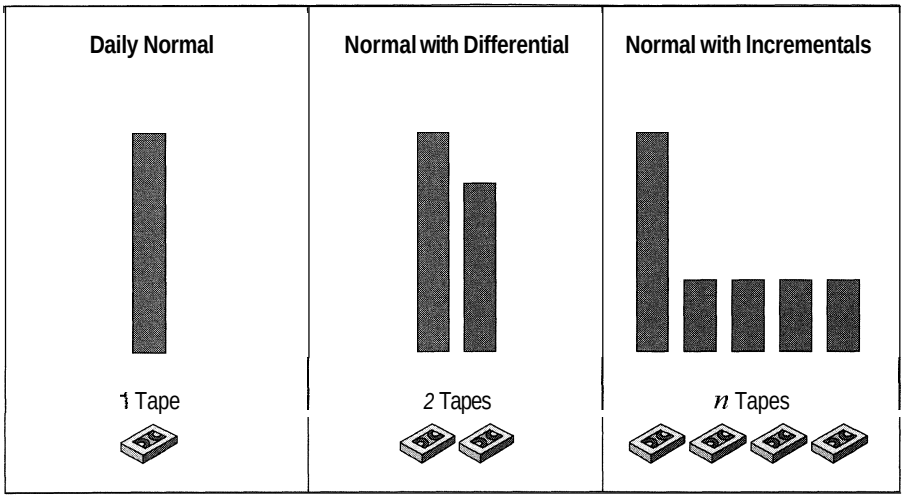


Figure 28.6 Restore strategies

Determining Overall Performance

The two factors that determine overall performance are the time needed to copy files from the tape to the hard drive and the time needed to replay transactions. If your system is properly configured, you should be able to reach the following speeds:

> **Note** These numbers will vary and you should perform tests on your systems.

- Backup to a single digital linear tape (DLT), approximately 30 GB per hour or 8.5 MB per second with hardware compression.
- Backup to a RAID 5 array of four DLT tape drives, approximately 40 GB per hour.
- Restoring to a RAID 5 disk partition, approximately 20 to 25 GB per hour with write-back caching.

  > **Note** If write-back caching is disabled, it takes twice as long to restore. Write-back caching should only be enabled if it is battery-backed.

Replaying transactions can take several hours on a large server because the database must run through every transaction since the last backup. The type of transaction that must be replayed also influences speed. It takes Web Storage System longer to replay log files with lots of deletes or attachments. Transaction replay can take, depending on the log file, from a few seconds to 4 minutes per log file. Exchange 2000 plays log files significantly faster than earlier versions.

Restore Categories

The type of restore you need to perform varies depending on the amount of information you need to restore. The restore process is the same for each category; however, it becomes more comprehensive depending on the level of restore used. The restore categories discussed are:

- Restoring online backups
- Restoring offline backups
- Restoring log files
- Restoring a single mail box
- Restoring a single database
- Restoring server clusters
- Full server restores

Restoring Online Backups

The Web Storage System databases can be restored to different server names. All transaction log files must be in place before starting the database. You cannot restore a full backup, start the database, stop it, and then restore incremental backups.

If you must rebuild a computer then you need a valid Exchange Server object in Active Directory. Use System Manager from another computer in the organization to verify that the original name of the server you want to recover still appears in Active Directory.

Restoring Offline Backups

Restoring offline backups is more time-consuming and less reliable than restoring online backups. However, the following are advantages of using an offline backup:

- You can move a database file if the hardware needs to be replaced or upgraded. Database services stop during an offline backup, so no changes take place until you restart the database.
- You can use an offline backup if an online backup fails because you receive a -1018 error. In this event, treat your backup as a temporary backup, and take all necessary steps to correct the situation before it deteriorates further.

The following are disadvantages of using an offline backup:

- You must stop database services because offline backup copies files to tape or disk.
- The database is not physically checked, which allows damaged data to be copied.
- The ability to play log files created after the backup is not supported in this scenario.
- The chance of data loss is increased due to file manipulations.

Restoring a Single Mailbox

If a single mailbox is completely removed, successfully restoring it requires that you use third-party software or restore the entire database to a recovery server. To prevent losing mailboxes it is recommended that you set retention times. To do this, open the **Mailbox Store Properties** dialog box in Exchange 2000. On the **Limits** tab, set the retention times. Under **Deletion** settings, type values for **Keep deleted items for __ days** and **Keep deleted mailboxes for __ days.**

If you delete a user account by accident, only the user account is removed. The mailbox itself is retained for 30 days, during which the user can reconnect it.

For mailboxes that are important, you might choose to have only that mailbox, or a select few mailboxes, in a single mailbox store.

If the user account is accidentally deleted, but the mailbox still exists, reconnect the mailbox. If a user's account is deleted, his or her mailbox is marked with a red x in the **Details** pane of Exchange System Manager.

To reconnect a deleted mailbox to a new user object

1. From Active Directory Users and Computers, create a new user object for the user.

   Important When creating the new user object, you must clear the **Create an Exchange Mailbox** check box. This is because you need to create a new Windows 2000 account without creating a corresponding Exchange mailbox. You connect this user account to a mailbox later.

2. From Exchange System Manager, navigate to the mailbox store where the user's mailbox is located.
3. In the **Details** pane, locate the mailbox for the user.

   Note Verify that the mailbox icon appears with a red x. Mailboxes that display with a red x are mailboxes that have been deleted, but have been retained in the mailbox store until the deleted mailbox retention period expires.

4. Right-click the mailbox with the user's name, and then click **Reconnect.**
5. In the **New User for this Mailbox** dialog box, select the new user object you created for the user in Active Directory Users and Computers, and then click **OK.**

Restoring a Single Database

If you need to restore a single database, dismount it from Web Storage System, and then use Windows 2000 Backup to restore the data. An extra storage group is dynamically created for the restore process and the database and transaction log files from the backup are restored in that process. The transaction logs are applied to the database in the restore storage group. After recovery, the consistent database is mounted into its original storage group by Web Storage System.

To restore a database to a different server, the database display name and the storage group display name must be the same. However, the organization name and administrative group name for the server to which you want to restore must match the server from which the database was backed up.

> **Note** The format of transaction log files is revised in Exchange 2000. When you upgrade from Exchange 5.5 to Exchange 2000, the existing transaction log files are removed and a new log series is created. Because of the log format change, you cannot restore an Exchange 5.5 database to an Exchange 2000 server.

Moving a Database to a Different Storage Group

When using Exchange 2000, you can move a database to a different storage group.

To move a database to a different storage group

1. Open the **Mailbox Store Properties** dialog box.
2. On the **Database** tab, select the **This database can be overwritten by a restore** check box.

This allows the database to be moved into a different storage group by a copy or a restore. The GUID in the database is changed to match the GUID in Active Directory.

Restore Constraints

Restoring data in a consistent way requires that you consider many different circumstances. Of particular interest to the network administrator is restoring Exchange data to a non-production server.

Restoring Web Storage System to a Different Server

You can restore Web Storage System to a Microsoft Exchange 2000 Server other than the one from which it was backed up. Use this method as a last resort to restore individual items, such as messages or folders, to a server other than the one in use, The secondary server must meet the hardware requirements to run Microsoft Exchange 2000 Server. It must not be connected to the network and must have enough disk space to restore the entire backup.

Do not use this method when a server fails because this method restores only Web Storage System; it does not restore the directory. When you finish restoring the backup. you must reconfigure permissions on the mailboxes.

Restoring Exchange Data to a Non-Production Server

You restore Exchange data to a non-production server to perform a single mailbox restore when the normal retention time has expired. You rarely need to perform this type of restore because Exchange 2000 has other features, such as Dumpster and mailbox retention, that make mailbox recovery simple. To restore data to a non-production server, you must fulfill the following requirements:

- The non-production restore server must reside on a separate Windows 2000 domain forest from the production network and servers.
- The storage group that hosts the restored database must have the same display name as the original production server.
- The database you want to restore must have the same display name as the original production server.
- The database name must be unique on the backup server in all storage groups. For example, if the database name is Priv.edb, there can only be one instance of a Priv.edb database on the secondary server.
- The organization name and the administrative group name must be the same.
- To recover the mailbox, reconnect the mailbox as described in "Restoring a Single Mailbox" earlier in this chapter. Use Outlook to connect the mailbox and download to the client computer.

Database Consistency

During normal operations of the ESE database, the database file is never completely up–to-date with changes made in memory. Although the database file and the memory pages make the ESE database consistent, the database files are inconsistent. Consistency of a database file is flagged in the database file header. If the database is stopped improperly, the flag indicates that when the database service restarts, the ESE will replay transactions and possibly rollback any incomplete transactions. When a database service shuts down properly, all transactions are written from memory to the transaction file. If no errors occur, the database closes in a consistent state and the flag is set to consistent.

If a database file is inconsistent, treat the contents of all database folders (mdbdata) as if they were a single file. Move them together and back them up together. Never move, rename, replace, or delete files in the database folders. Don't alter a file unless you know exactly why the database doesn't need it during startup. The following items should be considered when attempting to achieve database consistency:

- Defragmenting databases
- Soft recovery
- Log transaction redo
- Maintaining database integrity
- Using ESEUTIL
- Using ISINTEG

Defragmenting Databases

Defragmenting is rearranging data to fill database pages more efficiently. Defragmenting makes used storage contiguous and eliminates unused storage. There are two types of defragmenting: online and offline.

Online Defragmenting

Online defragmenting runs as part of database maintenance (this automatically detects and destroys objects that are no longer being used, thus freeing up database space) and does not decrease the size of the database. Online defragmenting is performed automatically at 2:00 AM every day by default. You can change this by altering settings in **Database maintenance,** on the **Exchange Database** tab.

Offline Defragmenting

Offline defragmentation creates a new database by copying all records and tables from the old database into the new database. Because this is a copy, defragmentation requires free disk space equal to the size of the database. After defragmentation is complete, ESE considers the new database to be a different database from the original. Therefore, the original database is deleted and its member log files cannot be replayed into the successor database. A full backup should be completed as soon as possible.

> **Caution** You should only defragment a database if a large number of users are moved off of a server, which frees up a large amount of memory. Offline defragmenting is not a recommended regular maintenance routine.

Database Size After Defragmentation

Although an ESE database can perform online defragmentation, it does not reduce the physical size of the database. At times, you can purge a large number of e-mail messages to free that memory. However, the database file size generally grows at a steady rate until the amount of data removed matches the amount of data added.

Offline defragmentation reduces the physical size of a database by writing data from the used pages to a new database created by the ESE Utility (ESEUTIL). This utility defragments, repairs, and checks the integrity of Exchange Web Storage System. When ESEUTIL starts, it creates a new database and copies database records to it. It discards unused pages. The result is a compact database file.

You should not perform defragmentation across the network. By default, ESEUTIL prefers the contents of a live database to a temporary one. When you use the command, TEMPDFRG.edb and TEMPDFRG.stm, the temporary new database replaces the original. There are switches to control where the temporary database is created and whether the original database file is deleted. It is not recommended to locate the temporary files across the network. This will cause defragmentation to take longer and will slow down the network during that time. For more information about ESEUTIL, see "Using ESEUSTIL" later in this chapter.

After defragmentation or repair, new database signatures are assigned. Because this is a new database, the database signature is changed; therefore previous transaction logs can no longer be replicated to this database.

Soft Recovery

In a soft recovery, a database starts normally and the storage group is initialized. Normally, the database file is in a consistent state, so the ESE simply begins to handle transactions. If the database is inconsistent—that is, it was not shut down properly—the ESE replays transactions from the checkpoint through to the log file, Edb.log. If the checkpoint file doesn't exist, the ESE starts with the earliest transaction log it finds and replays transactions. In this situation, you see events logged and the application event log shows log files being replayed. When the ESE finishes replaying the transaction, the database starts.

Log Transaction Redo

The ESE reads an operation, which includes the page that was changed. The ESE works through every operation in the transaction log files to the last log file, which might or might not be Edb.log. If the ESE finishes a series of operations and there is no commit operation for that transaction within the log file, the ESE reverses the other operations that make up the transaction. This process maintains the integrity of the database.

Using ESEUTIL

The ESEUTIL is a database-level utility that is not application specific. ESEUTIL extracts or dumps header information from database, checkpoint, and transaction log files to provide diagnostic information. The ability to dump file header information is especially useful when you want to check the current state of the database and transaction log files. You can use this information to discern why the database might not start, what transaction logs you need to replay, or to which log file the Edb.chk file points.

File Header Output

By dumping file header information, there are several output lines from which you can actually pull information. One is the DBSignature line, which is used to match the database file (*.edb), stream file (*.stm), and transaction logs. The signatures used by the ESE are created with a random number and the date stamp of when the signature was created. Another output line is State. This specifies whether a database is consistently shut down or not. A database restored from an online backup is always inconsistent, whereas a database restored from an offline backup is always consistent.

For example, the output of ESEUTIL/mh includes the following pieces of detailed information:

- DB Signature: Create time:03/30/2000 18:17:00 Rand:68947200 Computer
- cbDbPage: 4096
- State: Consistent
- Scrub Date: 00/00/1900 00:00:00
- Repair Date: 00/00/1900 00:00:00
- Last Consistent: (0x16,F,159) 0313112000 12:11:46
- Last Attach: (0x1,1404,17A) 03/31/2000 11:49:44
- Last Detach: (0x16,F,159) 03/3112000 12:11:46
- Log Signature: Create time:03/30/2000 18:16:56 Rand:68924870 Computer
- OS Version: (5.0.2195 SP 0)
- Previous Full Backup
- Current Incremental Backup

Note The term Last Consistent refers to the log file that was Edb.log at the last stop of the database. This gives the same three numbers as the Last Attach.

Transaction Log File Headers

To extract transaction log file header information, use ESEUTILlml and the transaction log. ESEUTILIml helps you determine why transaction logs don't play back into the database. For example, if you type ESEUTILIml edb00012.log, the output is the following parameters:

- *L-generation.* Log generation number in decimal.
- *Signature*. Transaction log file signature.
- *Database location.* Hard-coded into this file and cannot be changed except through the restore process.
- *Last-attach.* References the log file that was Edb.log the last time the ESE started.

These parameters produce three numbers. For example, if the number is: XXXYYYZZZ, *XXX* is a transactional log file number in decimal form, YYY is a page within the file, and ZZZ is the byte set within the page.

Repairing Damaged Databases

The repair function of ESEUTIL examines the structure of the database tables and records broken links. The ESE repair mode attempts to restore links, but the process is slow and data loss is highly likely. You should only use this as a last resort. Running the repair function ESEUTIL/p attempts to restore the damaged database. The output looks somewhat like ESEUTIL/g, but the database is read/write.

You might encounter physical damage to a database, where you receive -1018, -1019, or –1022 errors, the database page is removed, and data is lost. Any repair rewrites the database signature. You should also perform a full back up of the database. If you restore a backup made before the repair, then the database is rolled forward to the state it was in at the time of the repair. You cannot replay transaction logs that are created after the repair. You should never try to repair a directory database.

Checking Database Integrity

You can use ESEUTIL/g to check database integrity. This is a read-only utility that does not make changes to the database. It checks the database tables, rebuilds all indexes on a temporary database (Tempedb.edb), and compares the indexes. There are several switches to control the output and number of logs this command creates.

Forcing Soft Recovery

This unique property allows an administrator to deal with a potential disaster. For example, a server has three databases and experiences a hard drive failure. One of the databases was located on the lost hard drive. The server first attempts a soft recovery. When it encounters the lost database, it reports that the databases are no longer consistent and ceases the soft recovery. To get around this problem, the administrator can run ESEUTIL/R/I. R stands for soft recovery and

causes the server to resume the soft recovery. I stands for ignore hanging attach records, and causes the server to ignore the missing database. After transaction log redo is complete, the first two databases will be consistent. The missing database will have to be restored from backup.

ISINTE

Information Store Integrity Checker (ISINTEG) is a built-in component of Exchange 2000. ISINTEG is a suite of tests that check the database for inconsistencies. You should dismount the mail store or public store before you run ISINTEG. You should be selective about which tests you perform. Running the full complement of tests on a 40 GB database can take 12 hours to complete. Figure 28.7 shows the basic database schema.

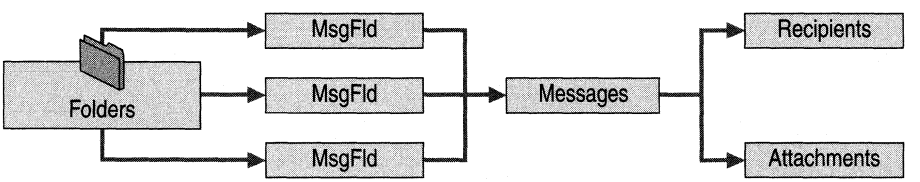


Figure 28.7 Basic database schema

Each Message Field (MsgFld) table connects a folder to its messages. It also contains per-recipient message properties and duplicates of message properties that have been requested for browsing. A MsgFld table is created for every folder in the Folders table the first time a message is created in that folder.

The Messages table contains all per-instance data for the message. It includes links to the Recipient tables and the Attachment tables. Messages can share recipient and attachment descriptions. There is one Messages table for the entire database.

The Attachments table contains all per-instance data for the attachment. There is one Attachments table for the entire database. The Attachments table is also a Message Folder table. It has a row in the Folders table, and the display name is MDB Attachments.

The Recipients table contains all per instance data for a recipient. There is one Recipient table for the entire database.

> **Note** If a storage group has more than five databases, ISINTEG will fail because it must be able to create a temporary database in order to run. Run ISINTEG with a maximum of five databases in the storage group. Customers must always dismount a database before running ISINTEG. If they have more than five databases in a storage group and they want to run ISINTEG, then in addition to dismounting the database they want to run against, they must also dismount the other databases so they have a maximum of five. ISINTEG can then create a temporary reference database and everything will work fine.

Reference Database and Multiple Passes

ISINTEG uses a reference database to check reference counts. It operates in three passes.

- **Pass 1: Constructing the reference database** In this pass, ISINTEG searches the entire original database and constructs a reference database that has many tables. Each table has as many rows as there are in the original database, but has only one column: the key. The key is copied from the original database. In ISINTEG's output, the test is labeled Reference Table Construction.
- **Pass 2: Filling the reference database** In this pass, the reference database is filled when the different tables in the original database are searched. This can be done in more than one test In this pass, random searches occur in the reference database tables.
- **Pass 3: Crosschecking the original database columns with the reference database columns** In this pass, every table in the reference database is searched sequentially, along with the corresponding table in the original database. Each column in the reference database is compared with the corresponding column in the original database and any discrepancies are reported (and corrected if in the -fix mode). In ISINTEG's output, the test is labeled Reference Table Verification.

Running ISINTEG

To run ISINTEG, type the following at the command prompt:

```
isinteg -s [-fix] [-verbose] [-l logfilename] [-t refdblocationl -test
testname[[, testname]...]
```

Table 28.6 shows the meaning of the terms you use to run ISINTEG.

Table 28.6 ISINTEG terms

Term Used on Command Line	Meaning
	Server name.
-fix	Specifies fix mode. ISINTEG will fix the problems it detects. The default is check-only mode, where problems are reported but are not fixed.
-verbose	Report verbosely.
-1 logfilename	Specifies the name of the output log file. By default, the results of check/fix operations are dumped to the Isinteg.pri or Isinteg.pub file, respectively, in the directory from which ISINTEG was executed.
-t refdblocation	Location of the temporary database (the reference database) that ISINTEG constructs during execution. This will be removed automatically by ISTNTEG when it terminates. Specifying the location of the temporary database on a different disk may improve performance because it facilitates parallel access.
-test testname	The ISINTEG test selected.

Recovering from Disasters

In this document, a disaster is defined as having to restore Exchange Server and/or Windows 2000 Server. With Exchange 2000 the introduction of multiple storage groups and databases adds more complexity to restoring. But in addition to using Windows 2000 Backup to back up and restore, you can reinstall Exchange 2000 using the /DisasterRecovery option. This option allows you to run Setup in Disaster Recovery mode to rebuild a server previously lost in the Exchange topology.

Besides Disaster Recovery setup, there are other procedures needed for recovering servers. In Exchange 2000, servers take on different roles such as Key Management Server and SRS. In addition to these new server roles, the platform Exchange is running on also adds more steps. Recovering an Exchange 2000 cluster server requires more steps than recovering a single Exchange 2000 member server.

This section describes recovery requirements and steps in the following scenarios:

- Recovering Exchange 2000
- Recovering an Exchange 2000 member server
- Recovering an Exchange 2000 member server running SRS
- Recovering an Exchange 2000 member server running Key Management Service
- Recovering an Exchange 2000 Cluster service

Requirements for Recovering E 2000

There are five common requirements for recovering all Exchange 2000 servers:

- **Windows 2000 and Exchange 2000 installation CDs** This includes any applicable service packs or hot fixes as outlined in the Windows 2000 and Exchange 2000 release notes.
- **Full backups of the system drive** This includes any other logical drives where critical application data is installed.
- **Recent Windows 2000 system state backup** A system state backup is an important type of Windows 2000 backup. Windows 2000 Backup captures and saves system configuration information that a file system backup normally fails to back up. System configuration information, such as the Windows registry and IIS metabase data, is included in a system state backup.
- **Backups of Exchange databases** These are online backups of the Web Storage System databases.
- **A server object in Active Directory for the server you want to restore** If an Exchange 2000 server is damaged, it is important that the server object in Active Directory is intact. If the server object is deleted from Active Directory, using either Exchange System Manager snap-in or Active Directory Users and Computers snap-in, recovery cannot succeed. Following a disaster, you must avoid changing the server object until the actual server is recovered.

  Important These steps assume that your Exchange 2000 server is a member server in a domain and not the server that runs Active Directory.

If the Exchange 2000 server is also an Active Directory server be sure to include backups of Active Directory in the system state backup. For information about the requirements for restoring Active Directory on a domain controller, see the *Windows* 2000 Sewer *Operations Guide.*

Recovering an Exchange 2000 Member Server

If Active Directory runs on a separate domain controller in the domain and it is intact, proceed with the steps below. If Active Directory is running on the same computer as the Exchange 2000 member server, you must first restore Active Directory before restoring Exchange 2000. For more information about restoring Active Directory, see the *Windows 2000 Sewer Operations Guide.*

In addition to reinstalling Windows 2000 on the computer and restoring file system backups, it is recommended that you restore the Windows 2000 system state. When you complete a backup of the Windows 2000 system state, which includes the Windows registry and IIS metabase, the computer returns to the state it was in before the backup.

Reinstalling Windows 2000

You might need to reinstall Windows 2000.

To reinstall Windows 2000

1. Install Windows 2000 as a stand-alone server.
2. Install the same version of Windows 2000 that was previously installed. For example, Advanced Server or Data Center Server.
3. Install Windows 2000 to the same hard drive and paths to which it was previously installed. If you need to, configure the drives to match the previous logical drive configuration.
4. Use the same server name as the original server.
5. Select all of the same components installed on the original server.
6. Do not rejoin the domain. Leave the server in a workgroup so that after you install Windows 2000, you can restore the system state that places the server in the correct domain.

   Important After reinstalling the correct version of Windows 2000, reinstall any service packs or hot fixes that were previously installed.

Restoring the System Drive

You should restore full backups of the system drive or any other logical drives where critical application data was installed. Use Windows 2000 Backup to restore file system backups to the computer on which you restored Windows 2000.

Restoring Windows 2000 System State

When you restore the Windows 2000 system state, the restored computer returns to its original domain where its computer account matches the System ID (SID) in Active Directory. Use Windows 2000 Backup to restore the Windows 2000 system state. After the restore, Windows 2000 Backup prompts you for the computer you want to restart.

> **Important** After restoring the system state and restarting the server, you might see an error message indicating that one or more services cannot be started. This includes services such as SMTP that were running prior to restoring system state, and services that have not been installed yet. These services require that you install Exchange 2000. If the full file system restore does not include the Exchange 2000 installation directory or other critical program data, this error message occurs. These services start after Exchange 2000 installs in disaster recovery mode.

Following the restore of your system state, the event log might show that some Exchange 2000 services have failed. If these services are not installed yet, when you restore the system state, Windows 2000 accepts that these services are installed on your server. These services start after Exchange 2000 installs in disaster recovery mode.

Run Exchange Setup in Disaster Recovery Mode

When you run Setup.exe with the /DisasterRecovery option, Exchange 2000 restores executable files and system settings without disturbing the existing Active Directory information for the system. Setup in disaster recovery mode installs Exchange 2000 without resetting the server's configuration to defaults, but instead, leaves the server in its last configuration.

To run Exchange Setup in disaster recovery mode

1. From your Exchange 2000 CD, run *.setup /DisasterRecovery*.
2. In Exchange 2000 Setup Wizard, be sure every component originally installed on the computer is set to the Disaster Recovery option. If the originally installed components are not selected for Disaster Recovery, select them manually.

   > **Important** You must install Exchange 2000 to the same drive and directory on which it was installed on the original server.

3. During disaster recovery, a message informs you that you cannot restore Exchange 2000 unless Active Directory still contains a server object for the server being restored. Use Exchange System Manager to verify that the server object exists for the server you are restoring. If the server object does not exist, the recovery process does not succeed.
4. Near the end of the setup process in disaster recovery node, you are prompted to restore databases and then restart. If setup finishes and another dialog box appears that prompts you to restart, ignore this message and restore the databases before you restart.

Recovering Databases

To recover databases in Web Storage System, you must verify that all services on which Exchange 2000 depends are running. To restore a database, you must also dismount the database. However, because Exchange supports multiple storage groups, you only need to dismount the specific database you want to restore. This allows users access to any other databases in Web Storage System.

To recover databases in Web Storage System

1. Use Windows 2000 Backup to restore your databases. In the **Restoring Database Store** dialog box, in **Temporary location,** specify a directory in which to store a log file that is different from the directory where the original log files exist. Make sure the location has enough disk space to store the files. If you restore databases or log files to their original location, any existing databases or log files are overwritten.
2. If you are restoring a full backup without any incremental backups, click **Last restore set** to start the log file after restoring the database. If you are restoring a backup with incremental backups, do not select this option until you restore the last incremental backup.
3. After you finish the restore, verify that the databases are mounted before you restart the system. The Exchange 2000 member server restores after you restart the system. Configuration settings that existed before the original server was damaged remain. The restored server can return to its previous role in the Exchange organization.

Recovering an Exchange 2000 Member Server Running Site Replication Service

Recovering an Exchange 2000 member server also running Site Replication Service (SRS) requires the same steps involved in recovering a single member server. However, there are additional steps to recover the SRS database after you run Disaster Recovery setup.

In addition to the requirements described in "Requirements for Recovering Exchange 2000" earlier in this chapter, back up the Exchange 2000 SRS database. Follow the same steps listed in "Recovering an Exchange 2000 Member Server" earlier in this chapter. Reinstall Windows 2000, restore the system drive backup, restore the Windows 2000 system state backup, run Exchange 2000 Setup in disaster recovery mode, and restore Exchange Web Storage System databases. However, prior to restarting the computer as described in the last steps listed in the "Recovering Databases" section of "Recovering an Exchange 2000 Member Server" earlier in this chapter, you must perform the steps listed below to recover the SRS database from backup. Recover the SRS database using Windows 2000 Backup, and then restart the server.

To enable and start SRS after disaster recovery

1. Using the Computer Management snap-in, under **Services and Application,** click **Services.**
2. Select **Exchange Site Replication Service** from the list of services, and then click **Properties.**
3. For the **Exchange Site Replication Service,** set the startup to **automatic,** and then start the service.

To reset the password for the Exchange 5.5 service account

1. Using the Exchange System Management snap-in, click **Administrative Groups,** and then select your site name.
2. Right-click the site name, and then click **Modify.** Clear the password and confirm password boxes. Re-type and confirm the password.
3. Click **Apply.**

Restoring an SRS Database

Using Windows 2000 Backup, follow the same steps listed in "Recovering an Exchange 2000 Member Server" earlier in this chapter. However, instead of selecting **Exchange Information Stores with Ntbackup,** select the SRS database to be restored.

When you restore the SRS database, you have completed the recovery of the Exchange 2000 member server running SRS. Before you restart the system, verify that you have performed all the steps described in "Recovering an Exchange 2000 Member Server" earlier in this chapter.

Recovering an Exchange 2000 Member Server Running Key Management Service

Recovering an Exchange 2000 member server also running Key Management Service requires the same steps involved in recovering a single member server. However, there are additional steps to recover the Certificate Authority (if Certificate Authority was running on the same server as Key Management Service) and Key Management Service database after you run Setup in disaster recovery mode.

In addition to the requirements described in "Requirements for Recovering Exchange 2000 earlier in this chapter, make sure you have the following backups:

- Backup of the Exchange 2000 Key Management Service database
- Backup of the system state, including Certificate Authority if the Key Management Service server also runs Certificate Authority

Follow the same steps listed in "Recovering an Exchange 2000 Member Server" earlier in this chapter. Reinstall Windows 2000, restore the system drive backup, restore the Windows 2000 system state backup, run Exchange 2000 Setup in disaster recovery mode, and then restore Exchange databases.

If the recovered server was also the Certificate Authority, in addition to being the Key Management Service; then by restoring the system state of the original server, Certificate Authority is also restored. Certificate Authority does not have to be checked as a component to install if you are restoring the system state.

Prior to restarting the computer as described in the last steps of "Recovering Databases" in the "Recovering an Exchange 2000 Member Server" section earlier in this chapter, perform these additional steps in recovering the Key Management Service database from backup. Then restore the Key Management Service database using Windows 2000 Backup and reboot the server.

To start Key Management Service following disaster recovery

1. Using the Computer Management snap-in, under **Services and Application,** click **Services.**
2. Select **Exchange Key Management Service** from the list of services, and then click **Properties.**
3. For the **Key Management Service,** type the Key Management password in **Startup Parameters,** and then select **Start.**

Restoring Key Management Service Database

Use Windows 2000 Backup to perform the steps as described in the "Recovering Databases" section in "Recovering an Exchange 2000 Member Server" earlier in this chapter. However, instead of selecting **Exchange Information Stores with NTBackup**, select the Key Management Service database you want to restore.

The Exchange 2000 member server running Key Management Service recovery is complete after you restore the Key Management Service database. Before you restart the system, verify that you completed the steps described in "Recovering an Exchange 2000 Member Server" earlier in this chapter.

Recovering an Exchange 2000 Cluster Server

Clustering provides a mechanism for moving resources between cluster nodes when a disaster occurs. In the case where a single node fails, clustering moves Exchange 2000 resources to another node in the cluster so that services remain available to users. The node that failed can be removed from the cluster and then later replaced with another node joining the cluster. Exchange resources can then be moved back to the newly joined node so that load balancing is again achieved. This section lists the steps involved in removing a non-functioning node from a cluster, rebuilding, and then rejoining the node to that cluster.

In addition to disasters involving the loss of a single node in a cluster, there are cases where the cluster-shared disk is lost. This section describes how to recover when the cluster quorum is lost.

Recovering a Single Server in a Cluster

Clustering provides recovery when a server node goes down in a cluster. When a single node in a cluster fails, Exchange resources running on the node are moved to an available node in the cluster. Exchange databases remain intact on shared storage and can be accessed by the Exchange virtual server from another node in the cluster. This provides reliability when disaster occurs on a single node in the cluster. Once resources are moved to an available node in the cluster, follow these procedures for removing the non-functioning node and replacing it with a new node.

Evicting the Lost Server Node from the Cluster

When one cluster node suffers a disaster and needs to be replaced by a new node, you must evict the lost node.

To evict the lost server node from the cluster

1. Use Cluster Administrator to remove the lost node from the cluster.
2. Use Cluster Administrator to verify that the evicted node for each cluster group and resource does not appear as a possible or preferred owner.
3. Physically remove the damaged node from the cluster and shared storage.

Building a New Server Node for the Cluster

The lost node does not have to be rebuilt like the original lost node. An entirely new node can be built (new computer name, new IP) and then joined to the cluster. Perform the following steps to build a new server node.

To build a new server node for the cluster

1. Install Windows 2000 on the new computer and provide a new computer name.
2. Join the same domain as before with same administrative permissions given to the Exchange Administrator account.
3. Set up the new computer to access the same shared storage as the original node.

Rejoining the Sewer Node to the Cluster

To rejoin the server node to the cluster, set up Cluster service on the newly built server. When asked to join a cluster, specify the cluster you want to join.

Installing Exchange on the Server Node and Moving Resources Back to the Node

You might need to install Exchange on the server node and move resources back to the node.

To install Exchange on the server node and move resources back to the node

1. Install Exchange 2000 on the newly joined node. You must do this before Exchange resources can be movfed back to the newly joined node.
2. Verify that the cluster groups and resources on the other node show that the new node is a possible or preferred owner.
3. Move the Exchange resources that originally failed back to the new node.

Recovering a Lost Cluster Quorum

In order to recover from a cluster quorum failure, you must perform a cluster quorum backup. In addition to the requirements described in "Requirements for Recovering Exchange 2000" earlier in this chapter, you must also have a system state backup containing the cluster quorum.

Typically you need to recover a single lost node in a cluster. However, you might have a case where the cluster quorum is lost on the shared disk along with the Exchange databases. If this occurs, you must restore the cluster quorum from backup.

Restoring the Cluster Quorum From Backup

Before you can restart Cluster service on any nodes in the cluster, you must restore the quorum.

To restore the quorum from backup

1. Use the *Windows* 2000 *Server Resource Kit* Dumpconfig utility to restore the signature of the quorum disk if it has changed since you made the backup.
2. If the Cluster service is running, stop the Cluster service on all cluster nodes.
3. Restore the system state (containing the contents of the cluster quorum disk) using Windows 2000 Backup. Windows 2000 Backup puts the contents of the cluster quorum disk in subdirectory *systemroot*\cluster\cluster_backup.
4. After restoring, you are prompted to restart. Instead of restarting, run Clusrest.exe tool to restore the content of the *systemroot*\cluster\cluster_backup directory to the cluster quorum disk. The Clusrest.exe tool is included in the *Windows* 2000 *Server Resource* ***Kit.***
5. Restart the computer.

Restoring Exchange 2000 Databases from Backup

After you restore the quorum and restart the nodes in the cluster, verify whether the shared disk resource can be accessed after the Cluster service has started. If the shared disk where Exchange databases reside can be accessed and has survived the disaster, check to see if the .edb, .stm, and .log files still exist for the Exchange virtual server storage groups. If they are intact, start your Exchange resources. If the shared drive is lost, restore your Exchange databases from backup.

To restore Exchange 2000 databases from backup

1. Start Exchange System Manager, and then select the **Do not mount at startup** check box for databases owned by the Exchange virtual servers on the cluster. This avoids creating new databases on the shared disk resource when the Exchange resources start.
2. Using Windows 2000 Backup, perform the steps as described in the "Recovering Databases" section in "Recovering an Exchange 2000 Member Server" earlier in this chapter.

   Note On a cluster server, you must verify that the shares where Exchange databases reside are available to and accessible by the cluster node that owns the disk resource.

3. Use Exchange System Manager to verify that databases are mounted and check the Event log. Click to clear the **Do not mount at startup** check box for each database that is successfully restored.

Best Practices

There are many recommended procedures for achieving successful backups and restores. The following sections outline some of the best practices for specific backup and restore situations.

Increase Backup Tape Reliability

Note the capacity of you backup tapes in relation to the size of the database. Ensure that the raw storage capacity of your tape exceeds the compressed storage capacity of your database by a comfortable safety margin. If it does not, plan for tape changes when doing backups. For example, you must change a tape with a raw capacity of 4 GB twice if you use it to back up an uncompressed database of 10 GBs, assuming the data is compressed during the backup. You must be careful because all DLT drives use the LZW (Lempel-Ziv-Welsh) compression algorithm. Compressed data or random data (such as encrypted data) can actually expand by about five percent if you attempt to compress it further, which can reduce drive throughput.

Note The LZW compression algorithm is an industry standard that reduces the number of bits to transfer without losing data. Transmitting data is compressed at a ratio of 2.8:1 at 9.6 kilobytes per second (Kbps). This is equivalent to transmitting non-compressed data at 27 Kbps. Such numbers reflect an ideal situation, and a ratio of 1.5-1.8:1 is the normal experience. The switches with this feature can transmit in either compressed or non-compressed mode.

To increase reliability, take the following precautions:

- Routinely clean the tape drives according to manufacturer specifications.
- Do not overuse tapes. Discard them after they reach the maximum number of cycles specified by the manufacturer.
- Store the tapes in a safe and accessible location.

Verify and Validate Backups

The ability to restore data and servers depends on the quality of your backups. Therefore. it is important to verify the success of your backup procedure. For complete fault tolerance, perform verification at two levels: verify the event and verify the data.

> **Note** If you turn on detailed logging for Windows 2000 Backup, the reports contain more detailed data about what is actually backed up.

Verifying the Event

It is important to verify that the backup occurs without errors. Examine the Windows 2000 Backup log for the backup event you want to verify. Thoroughly review the events in the Windows 2000 event log to ensure your backup completed as scheduled without errors. You should research and resolve errors or inconsistencies in the logs as soon as possible.

Verifying Data

Verifying data includes checking the data from the storage device to another computer to verify that the backup works correctly. Verifying all backups from all servers in a large installation is difficult. However, by rotating a restore process to include various servers or backup devices, you can verify the integrity of the system and identify potential problems before you lose data. This process also helps to train new administrators to perform restore procedures.

Document and Archive Backups

It is important to document your backup strategy and provide step-by-step instructions that describe how to use backups to restore data. Backup and disaster restore procedures should also specify escalation procedures and the members of the organization responsible for approving the escalation of the situation if any part of the process fails. Label backup media and store them in a secure location. For maximum reliability, archive full backups at a separate location, preferably offsite.

Test Backups

To prepare for server failures, practice your backups using different scenarios. It is much easier to deal with an emergency when you are prepared. Run your backups on a test server, duplicating your working conditions on a weekly or monthly basis until you are proficient. Be sure that you can deal with the following three scenarios:

- **Full restore of Windows 2000** This can happen if a Service Pack or patch is not installed correctly.
- **Full restore of Microsoft Exchange 2000** This can happen if there is system or database damage.
- **Full restore of Microsoft Exchange 2000 and Windows 2000** This can happen if there is a power outage.

  Note Do not assume that a database is damaged. The only error that can be relied on to prove damage is the repeatable error message -1018. If you assume incorrectly that a database is damaged this might lead you to take drastic measures that can prolong downtime and cause unnecessary data loss. Most startup failures that are blamed on damage are really caused by the presence of mismatched files or the absence of necessary files.

CHAPTER 29

Monitoring and Maintaining

Monitoring system activity and server performance, when combined with disaster planning and regular backup, is a necessary part of preventive maintenance for the server running Microsoft Exchange 2000 Server. Through monitoring, you obtain data that you can use to diagnose system problems, plan growth, and troubleshoot problems. You can use the Exchange 2000 Monitoring and Status tool, diagnostic logging, extended logging, and Queue Viewer, to keep up-to-date on the status of your Exchange 2000 servers. You can also use Microsoft Windows 2000–based tools such as Performance Monitor, Event Viewer, Task Manager, and Terminal Services Client, to ensure you have current information about how Exchange Server and the network are operating. Two additional services, Network Monitor and the DOS-based Network Diagnosis tool (Netdiag), provide additional network monitoring information.

Note Because of architectural differences between the monitoring user interfaces in Microsoft Exchange 5.5 and Exchange 2000 Server, mixed mode networks require that you use both monitoring systems. You can use the Exchange 2000 Server monitoring system to monitor Exchange 2000 Server only, and you can use the Exchange 5.5 monitoring system to monitor Exchange 5.5 only. However, when you are in native mode you can monitor by using only the Exchange 2000 Server monitoring system.

In This Chapter

Exchange 2000 Monitoring Features

Windows 2000 Performance Monitoring Tools

Analyzing Performance Data

Exchange 2000 Monitoring Features

Exchange 2000 provides a number of features that assist you in monitoring and maintaining your Exchange server and network. These tools include the Monitoring and Status tool, extended logging and diagnostic logging.

Monitoring and Status Tool

The Monitoring and Status tool, located in the Tools folder in Exchange System Manager, is the primary Exchange 2000 tool you use to monitor the health and status of your network and servers. The tool is composed of two user interfaces: Notifications and Status.

Notifications

You can use the Notifications user interface with the Status user interface to set up e-mail alerts or script triggers when a warning or critical state is reached on any of your network servers. In the Status user interface, you can configure warning and critical states for Simple Mail Transfer Protocol (SMTP) and X.400 queues, available virtual memory, CPU activity, and free hard disk space for an array of Microsoft Windows NT 4.0, Windows 2000, and Exchange 2000 services.

Configuring Warning and Critical States

You can configure Exchange to constantly monitor the performance levels of an array of network and application services. Levels for both warning states and critical states can be established so problems are announced and can be dealt with as they occur. You configure these using the Status and Notifications user interface.

Effective monitoring requires that you establish levels of acceptable performance for each resource. You can determine a warning threshold and a critical state threshold from this baseline level.

Exchange 2000 logs a critical state when any of the following default Microsoft Exchange Services stops running:

- Web Storage System
- Message Transfer Agent (MTA) Stacks
- Routing Engine
- System Attendant
- SMTP
- World Wide Web Publishing Service

To add a service to the default Microsoft Exchange Services

1. In the physical server's **Properties** dialog box, click **Detail** on the **Monitoring** tab.
2. Select a service in the **Default Microsoft Exchange Services** dialog box.
3. Click **Add,** select the service you want to add, and then click OK

You can add other services by performing the following steps.

To add a resource

1. In the physical server's **Properties** dialog box, click **Add** on the **Monitoring** tab.
2. Select a resource from the list, and then click **OK.**
3. Configure the resource with the **Warning state** and **Critical state** thresholds in the **Thresholds** dialog box of the resource, and then click **OK.**

> **Note** Any service that you add to the default Microsoft Exchange Services follows the same configuration rules as the default set. This means a critical error occurs if the selected service stops running. To add a set of services so a warning message rather than a critical error message occurs, you must create another set.

Note You use most of these services exclusively for troubleshooting. Many services start and stop frequently in normal operating conditions; therefore, you must think carefully before configuring Exchange 2000 to generate errors when a service stops running. Also, because monitoring and generating notifications consumes server resources, it is recommended that you do not configure Exchange 2000 to monitor an excessive number of services.

Configuring Notifications

Because you cannot instantly access all elements of every network server, Exchange provides a Notifications tool that you can set to trigger a script or to alert appropriate personnel when Exchange crosses a warning or critical threshold.

> **Note** Exchange 2000 is not configured by default to send notifications; all notifications must be configured through the user interface.

For example, you can set a five-minute warning threshold and a 10-minute critical threshold for SMTP queue growth, and configure Exchange to notify you through e-mail. When the SMTP queue growth reaches five minutes, Exchange sends the following message by default:

```
%TargetInstance.Name% has reported a %TargetInstance.ServerStateString%.
Reported status is:
Queues - %TargetInstance.QueuesStateString%
Drives - %TargetInstance.DisksStateString%
Services - %TargetInstance.ServicesStateString%
```

> **Note** *%TargetInstance.Name%, %TargetInstance.ServerStateString%,* and other fields are placeholders in which Exchange places data about the specifics of the warning, including the server name and instance data.

For more information about setting Notifications, see Exchange 2000 Server Help.

Status

The Status user interface allows you to view servers and connectors on your network with their status condition and administrative group designation. You can also disable all monitoring of a server through this user interface.

Monitoring the status of servers and connectors is one of the most effective ways to ensure that your network functions correctly. The status window of the Monitoring and Status tool shows the operating status of each server and connector on the network.

The following status states and definitions apply to servers:

- Available: Server online and functioning normally.
- Unreachable: One of the primary services on the server is down.

  Note If a server is unreachable and is in a different routing group, it may indicate that a connector between routing groups is down or does not exist.
- In Maintenance Mode: Monitoring is disabled on this server for maintenance, backup, repair, or another reason.
- Unknown: System Attendant cannot communicate with the local server.

The following status states and definitions apply to connectors:

- Available: Functioning properly.
- Unavailable: A communication function such as routing service is not functioning on this connector.

When you want to disable monitoring on a server for maintenance, backup, or recovery, select the **Disable all monitoring on this server** check box on the **Monitoring** tab of the **Server Properties** dialog box for an individual physical server.

Diagnostic Logging

You can use diagnostic logging in Exchange 2000 to monitor protocol connectors and Microsoft Exchange connectors. This is an effective way to keep up-to-date on a server's status and to prevent potential problems. Diagnostic logging produces data about monitored connectors, which you can view using Event Viewer, one of the Windows 2000 Administrative Tools.

Exchange allows you to log diagnostic data by using the physical server's **Properties** menu. On the **Diagnostic Logging** tab, select a service to monitor from the **Services** pane and one or more categories to monitor from the **Categories** pane.

Table 29.1 Services available for diagnostic logging

Service to Monitor	Abbreviation to Select
Internet Message Access Protocol Version 4 (IMAP4)	IMAP4Svc
Lotus Notes GroupWise	LME-GWISE
Lotus Notes	LME-NOTES
Address List	MSExchangeAL
Lotus cc:Mail	MSExchangeCCMC
Microsoft Exchange Directory Synchronization	MSExchangeDX
Microsoft Exchange Schedule Plus Free/Busy	MSExchangeFB
Microsoft Exchange Router for Novell GroupWise	MSExchangeGWRtr
Information Store System	MSExchangeIS\System
Information Store Mailbox	MSExchangeIS\Private
Information Store Public Folders	MSExchangeIS\Public
Microsoft Mail	MSExchangeMSMI
MAPI Address Book Proxy Service	MSExchange NSPI Proxy
MAPI Address Book Referral Service	MSExchangeRFR Interface
Site Replication Service	MSExchangeSRS
SMTP Routing Engine and Transport	MSExchangeTransport
POP3 Protocol	POP3Svc

You can configure logging levels for most services and categories. When a connector tries to send a log message, it first checks the logging level of the message against the logging level that is configured. If the logging level of the message is the same or higher than the logging level configured, the message is logged. Otherwise, the connector does not log a message and continues.

Logging level options are as follows:

- None: Only error messages are logged.
- Minimum: Warning messages and error messages are logged.
- Medium: Informational messages, warning messages, and error messages are logged
- Maximum: Troubleshooting messages (fine and detailed information), informational messages, warning messages, and error messages are logged.

Note It is not recommended that you use the maximum logging setting unless instructed to do so by Microsoft Technical Support because maximum logging considerably drains resources.

Protocol Logging Tool

The Protocol Logging tool provides detailed information about the commands being sent and received by SMTP and Network News Transfer Protocol (NNTP). This tool is particularly useful in monitoring and troubleshooting protocol or messaging errors.

The user interface for SMTP and NNTP logging is located in the **Properties** dialog box of an individual SMTP or NNTP virtual server.

To configure SMTP and NNTP logging

1. Select the **Protocols** folder, and then select either the **SMTP** or **NNTP** folder.
2. Right-click the virtual server. Exchange displays the **Default Virtual Server Properties** dialog box.
3. Select **Enable Logging,** and then click **Properties.**

Figure 29.1 shows the **General** tab in the **Extended Logging Properties** dialog box.

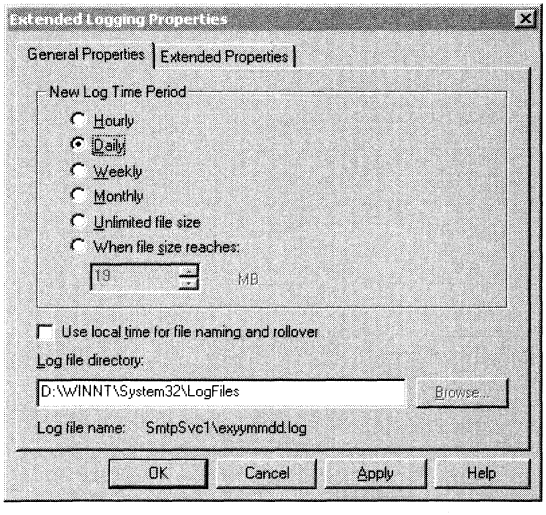


Figure 29.1 General tab in the Extended Logging Properties dialog box

You can select the logging frequency and the name and location of the log file to create. To do so, click **Properties** on the **General** tab of the **Default Virtual Server Properties** dialog box. By default, the log file path is:

Systemroot:\WINNT\System32\LogFiles

You can change the path or log file name by clicking one of the option buttons or by letting the current path stand as the default setting. You must select a logging format for the SMTP or NNTP log. By default, the log file name uses the following creation date for SMTP server:

SmtpSvc1\exyymmdd.log.

Figure 29.2 shows the **Extended Properties** tab in the **Extended Logging Properties** dialog box.

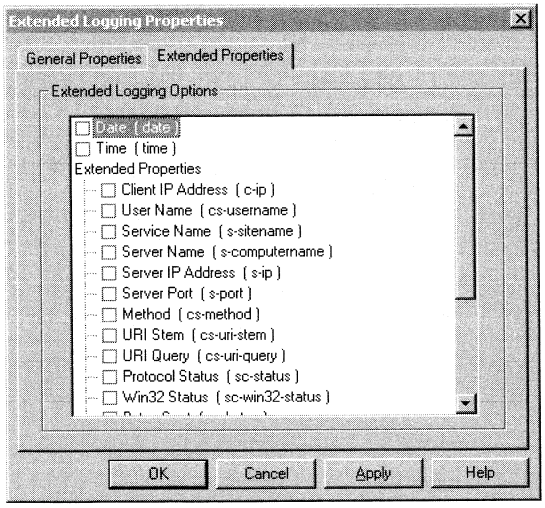


Figure 29.2 Extended Properties tab in the Extended Logging Properties dialog box

On the **Extended Properties** tab, you can select the configuration parameters from the **Extended Logging Options** pane. These logging options provide more detailed information on the Service and Category logging properties already configured for the virtual server. They do not add additional services or categories to be logged. Unlike the IMAP4 and POP3 protocol connectors and the MSExchange connectors, you cannot select the services and categories to log for SMTP or NNTP.

You can also establish a rollover time frame or file size using **New Log Time Period** on the **General Properties** tab in the **Extended Logging Properties** dialog box. When the time interval expires or when the log file reaches the set size, logs are overwritten.

Using Event Viewer to View Logs

Event Viewer is a Microsoft Management Console (MMC) snap-in that provides event information about running applications, the directory service, the file replication service, security, and the system. It also allows you to view the logs you configure for IMAP4 and POP3 protocol connectors and the MSExchange connectors. Events are logged by date, time, source, category,

event number, user, and computer. By viewing the event data, you can see errors and warnings, and diagnostic information to find the problems that occur on any computer in the network. You can see an event's logging properties and text by viewing the properties of the event. For more information about using Event Viewer, see Windows 2000 Help.

Queue Viewer

As part of the monitoring process, you can view the X.400 and SMTP queues, and other connectors that are installed on a server by using Exchange System Manager. You can use information such as message age and the number of messages in the queue to troubleshoot problems on a server.

You access the queues through three paths, as shown in Figure 29.3.

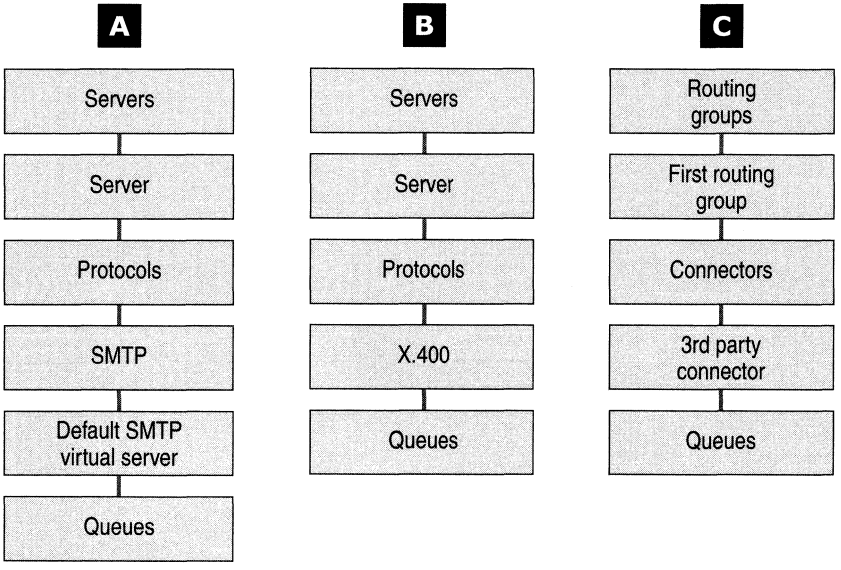


Figure 29.3 Paths to accessing queues

The two queues most useful to monitor are: Local Delivery and the Messages Awaiting Directory Lookup. A backlog in the Local Delivery queue indicates a problem with Web Storage System. A backlog in the Messages Awaiting Directory Lookup queue indicates there is a problem contacting the domain controller.

In addition, a backlog of messages with the same destination indicates there might be a problem with the destination domain controller.

Windows 2000 Performance Monitoring Tools

Exchange 2000 is integrally linked to Windows 2000, so many of the tools you use to monitor Exchange Server and the network are part of the Windows 2000 operating system. You can use the Windows 2000 tools Performance console, Netdiag, Task Manager, and Network Monitor to monitor performance.

Performance Console

System performance in Exchange is monitored in part by the Performance console, which includes System Monitor and Performance Logs and Alerts.

System Monitor

System Monitor displays graph, histogram, or report displays of system data. System Monitor provides short-term viewing of data, and information for troubleshooting and diagnosis. It includes tools such as physical hard disk counters and workload balance tools.

Performance objects that are associated with a resource or service that you can monitor contain at least one performance counter. You can view selected performance counters individually or in relation to other available counters.

> **Note** Monitoring large numbers of counters can create overhead, which can make the system unresponsive to keyboard or mouse input. To reduce this burden, you can either display data in report view when collecting information from a large numbers of counters, or direct data to a binary log and view the data in System Monitor as it is written to the log.

For more information about System Monitor, see Windows 2000 Help.

Performance Objects

When you monitor Exchange 2000 performance, you often rely on data contained in Windows 2000 performance objects, which is collected from components in your computer and monitored in Windows 2000 System Monitor and Performance Logs and Alerts. As a component functions in your system, it generates performance data. The data is formulated into performance objects that are typically named for the component generating the data. For example, the Processor object is a collection of performance data about processors on your system.

A range of performance objects is built into the Windows 2000 operating system and typically corresponds to the major hardware components such as memory, processors, and so on; other applications might install their own performance objects.

Exchange 2000 installs its own set of performance objects and counters. Table 29.2 provides information about the services or resources in Exchange that you can monitor using System Monitor.

Table 29.2 Exchange services or resources monitored using System Monitor

Service or Resource to Monitor	Performance Object
Active Directory DXA Connector	MSExchangeADDXA
Address List	MSExchangeAL
Chat Communities	MSExchange Chat Communities
Chat Service	MSExchange Chat Service
Directory Service Access Caches	MSExchangeDSAccess Caches
Directory Service Access Contexts	MSExchangeDSAccess Contexts
Directory Service Access Processes	MSExchangeDSAccess Processes
Document Conferences	MSExchangeCONF
Document Conferencing Manager	MSExchangeDcsMgr
Document Conferencing Protocol (Multipoint Control Unit)	MSExchangeT.120
Epoxy Queues and Activity	EXIPC
Event Store	MSExchangeES
File Replication Connector	FileReplicaConn
File Replication Settings	FileRepSet
HTTP Extension	Exchange Server HTTP Extension
Internet Information Server Store Driver	Exchange Store Driver (IIS)
Internet Message Access Protocol Version 4	MSExchangeIMAP4
Web Storage System	MSExchangeIS
Private Information Store	MSExchangeIS Mailbox
Public Information Store	MSExchangeIS Public
System Information Store	MSExchangeIS

Table 29.2 Exchange services or resources monitored using System Monitor *(continued)*

Service or Resource to Monitor	Performance Object
Mailbox Information Store	MSExchangeIS Mailbox
Public Folders Information Store	MSExchangeIS Public
Lotus CC Mail	MSExchangeCCMC
Lotus Notes Message Center	MSExchangeNMC
Message Transfer Agent	MSExchangeMTA
Message Transfer Agent Connections	MSExchangeMTA Connections
MS Mail Connector Interchange	MSExchangeMSMI
Exchange Referral Service	MSExchangeSA-RFR
MS Mail Connector Mail Transfer Agent	MSExchangePCMTA
Name Service Provider Interface (Active Directory Integration)	MSExchangeSA-NSPI Proxy
Network News Transfer Protocol Commands	NNTP Commands
Network News Transfer Protocol Server	NNTP Server
Novell Groupwise Connector	MSExchangeGWC
Object Linking and Embedding database events	MSExchangeOledb Events
Object Linking and Embedding database resources	MSExchangeOledb Resources
Post Office Protocol Version 3	MSExchangePOP3
Service Account	MSExchangeSA
Site Replication Service	MSExchangeSRS
Simple Mail Transfer Protocol	SMTP
Store Driver	Exchange Store Driver (Store)
Video Conferencing	MSExchangeIPConf
Web Mail	MSExchangeWebMail

Each Exchange 2000 performance object has at least one associated counter, which you can configure to monitor instances involving any number of server actions. For more information about the function of a counter for a specified performance object, click **Select Counters from List,** select a counter, and then click **Explain.**

Physical Hard Disk Counters

Statistics about hard disk drive usage help you balance the workload of network servers. System Monitor provides physical hard disk counters for troubleshooting, capacity planning, and measuring activity on a physical hard disk.

At a minimum you need to monitor the following counters:

- PhysicalDisk\ Disk Reads/sec and Disk Writes/sec
- PhysicalDisk\ Current Disk Queue Length
- PhysicalDisk\ % Disk Time
- LogicalDisk\ % Free Space

When testing hard disk performance, you can log performance data to another hard disk or computer so it does not interfere with the hard disk you are testing.

You might want to observe additional counters, such as:

- PhysicalDisk\ Avg. Disk sec/Transfer
- PhysicalDisk\ Avg. Disk Bytes/Transfer
- PhysicalDisk\ Disk Byteslsec

The PhysicalDisk\ Avg. Disk sec/Transfer counter reflects how much time a hard disk takes to fulfill requests. A high value may indicate the hard disk controller is continually retrying the hard disk because of failures. These failures increase average hard disk transfer time. For most hard disks, high average hard disk transfer times correspond to values greater than 0.3 seconds.

You can also check the value of PhysicalDisk\ Avg. Disk Bytes/Transfer. A value greater than 20 KB indicates the hard disk is performing well. Low values result when an application accesses a hard disk inefficiently. For example, applications that access a hard disk at random raise PhysicalDisk\ Avg. Disk sec/Transfer times because random transfers require increased seek time. PhysicalDisk\ Disk Byteslsec provides the throughput rate of your hard disk system.

Because hard disk counters can cause a modest increase in hard disk access time, Windows 2000 does not automatically activate the counters at system startup.

Logical Disk Counters

Performance data about the logical hard disk is not collected by the operating system by default. To obtain performance data for logical drives or storage volumes, you must type **diskperf –yv** at the command prompt. This causes the hard disk performance statistics driver, which is used for collecting hard disk performance data, to report data for logical drives or storage volumes. Windows 2000 uses the **diskperf –yd** command to obtain physical drive data. For more information about using the **diskperf** command, type **diskperf** –? at the command prompt.

Determining Workload Balance

To balance loads on network servers, you need to know how busy the server hard disks are. You can determine this by using the PhysicalDisk\ % Disk Time counter, which indicates the percentage of time a drive is active. If the results of PhysicalDisk\ % Disk Time is high (more than 90 percent), you can check the PhysicalDisk\ Current Disk Queue Length counter to see how many system requests are waiting for hard disk access. The waiting input/output (I/O) requests should be sustained at a number no more than 1.5 to 2 times the number of physical hard disk spindles.

Most hard disks have one spindle, although redundant array of independent disks (RAID) devices usually have more. A RAID device appears as one physical hard disk in System Monitor. RAID devices created through software appear as multiple drives (or instances). You can either monitor the physical hard disk counters for each physical drive (other than RAID), or you can click **All Instances** on the **Add Counters** dialog box to monitor data for all of the computer's drives.

Use the values of the PhysicalDisk\ Current Disk Queue Length and PhysicalDisk\ % Disk Time counters to detect bottlenecks in the hard disk subsystem. If the PhysicalDisk\ Current Disk Queue Length and PhysicalDisk\ % Disk Time values are consistently high, consider upgrading the hard disk or moving some files to an additional hard disk or server.

The system maps physical drives to logical drives using the same instance name. Therefore, if you have a dynamic volume that consists of multiple physical hard disks, instances might appear as **Disk 0 C:, Disk 1 C:,** and **Disk 2 D:,** (where C: is made up of physical drives 0 and 1). If you have two logical partitions on a hard disk, the instance appears as **0 C:** D:.

For hardware-enabled stripe sets, statistics for each hard disk are not available. You can obtain this data only when monitoring stripe sets enabled in software.

If you use a RAID device, the PhysicalDisk\ % Disk Time counter can indicate a value greater than 100 percent. If this happens, use the Avg. Disk Queue Length counter to determine the average number of system requests waiting for hard disk access.

Performance Logs and Alerts

Performance Logs and Alerts contain features for logging counter and event-tracing data, and for generating performance alerts. With counter logs, you can:

- Record data about hardware usage and the activity of system services from local or remote computers.
- Configure logging to occur manually on demand or automatically based on a user-defined schedule.
- Enable continuous logging, which is subject to file size or duration limits.
- View logged data by using the System Monitor display or by exporting it to a spreadsheet program or database so you can analyze the data and generate a report.

Trace logs record data when an event such as a hard disk I/O error or page fault occurs. When the event occurs, the monitoring service sends the data to the log service.

You can set an alert that sends a message, runs a program, or starts a log when a selected counter's value equals, exceeds, or falls below a specified setting.

For more information about Performance Logs and Alerts, see Windows 2000 Help.

When you select a tool to see its status, if any logs or alerts are defined, then they appear in the details pane. A sample settings file for a counter log named System Overview is included with Windows 2000 and appears in the right pane of the Performance console when you select **Counter Logs** under **Performance Logs and Alerts.** You can use this file to see basic system data such as memory activity, hard disk activity, and processor activity.

You can right-click the log icon to create a new log or alert in a new file, or you can use settings from an existing HTML file as a template.

Task Manager

Task Manager provides information about processes, memory usage, and processor performance statistics. However, it lacks the logging and alert capabilities of the Performance console and does not provide the breadth of information available from System Monitor counters.

Monitoring Processes

To use Task Manager to monitor a process, click the Task Manager's **Processes** tab to see a list of processes that are running and information about their performance. Task Manager process tables include all processes that run in their own address space, including all applications and system services. Exchange 2000 runs two processes you can view using Task Manager: Store.exe and Inetinfo.exe.

> **Note** Even if you have more than one virtual server running on a server, only one instance of Store.exe and Inetinfo.exe exists.

Monitoring the System

To use Task Manager to see a dynamic overview of system performance, including a graph and numeric display of processor and memory usage, click the **Performance** tab in the **Windows Task Manager** dialog box.

To graph the percentage of processor time in privileged or kernel mode, on the **View** menu, click **Show Kernel Times.** This is a measure of the time that applications take to operate system services. The remaining time, known as user mode, is spent running threads in the application code.

If you use multiprocessor computers, on the **View** menu, select **CPU History,** and then view the non-idle time of each processor in a single graph or in separate graphs.

For more information about Task Manager, see Windows 2000 Help.

Terminal Services Client

The Terminal Services Client user interface, located in Windows *2000* Administrative Tools, allows you to log on to a remote network computer from another terminal. You can use this user interface to perform remote administration and monitoring of any accessible server on the network.

For example, if you are in Seattle and you discover that a server in Tokyo is not responding and you cannot contact the on-site administrator, but the server has Terminal Services Client installed, you can remotely log on to the computer and administer it. You can restart it, run installed applications, start and stop services, and monitor performance.

Another benefit of Terminal Services Client is that you can run applications that are not installed on your computer by gaining access to a computer with Terminal Services Client that has the application you want to run.

For more information about Terminal Services Client, see Windows *2000* Support Tools Help.

Network Monitor

Monitoring a network typically involves observing resource usage on a server and measuring network traffic. You can use Network Monitor to understand the traffic and behavior of your network components. Unlike System Monitor, which you use to monitor hardware and software, Network Monitor exclusively monitors network activity. You can use System Monitor to monitor hardware and software on the network; however, for in-depth traffic analysis, you should use Network Monitor.

Observing Resource Usage

To check resource usage, start by tracking the counters on your server. To focus on network resource usage, monitor the counters that correspond to the various layers of your network configuration. Abnormal network counter values often indicate problems with a server's memory, processor, or hard disks. For this reason, the best approach to monitor a server is to watch network counters with the following performance counters: Processor\ % Processor Time, PhysicalDisk\ % Disk Time, and Memory\ Pages/sec.

For example, if a dramatic increase in pages per second is accompanied by a decrease in total bytes per second handled by a server, the computer probably lacks physical memory for network operations. Most network resources, including network adapters and protocol software, use non-paged memory. A computer can page excessively if most of its physical memory is allocated to network activities, which leaves a small amount of memory for processes that use paged memory. To verify this usage, check the computer's system event log for entries indicating it is out of paged or non-paged memory.

For more information about the counters, see "Physical Hard Disk Counters" earlier in this chapter.

Measuring Network Traffic

You can use Network Monitor to observe throughput across network layers. Investigating network performance includes monitoring activity at different network layers. There are four layers in which you monitor network activity: the data-link layer, the network layer, the transport layer, and the presentation or program layer.

Data-Link Layer

The data-link layer includes the network adapter. Use the following network user interface object counters to monitor network activity at the data-link layer:

- Bytes total/sec
- Bytes sent/sec
- Bytes received/sec

Network Layer

Use the Internet Protocol (IP) object counters to monitor network activity at the network layer:

- Datagrams Forwarded/sec
- Datagrams Received/sec
- Datagrams/sec
- Datagrams Sent/sec

Transport Layer

The transport layer varies with network protocol in use. For TCP/IP, use the TCP object counters to monitor network activity at the transport layer:

- Segments Received/sec
- Segments Retransmitted/sec
- Segments/sec
- Segments Sent/sec

If the retransmission rate is high, this may indicate a hardware problem.

The Internet Control Message Protocol (ICMP) and User Data Protocol (UDP) object counters are useful for more extensive monitoring of TCPIIP network transmissions. The ICMP performance object consists of counters that measure the rates at which ICMP messages are sent and received by using the ICMP protocol. It also includes counters that monitor ICMP protocol errors. The UDP performance object consists of counters that measure the rates at which UDP datagrams are sent and received using the UDP. It includes counters that monitor UDP errors.

If you use the NetBIOS Enhanced User Interface (NetBEUI) protocol, use the following counters:

- NetBEUI\ Frame Bytes Receivedlsec
- NetBEUI\ Frames Receivedlsec
- NetBEUI\ Frames Rejected/sec
- NetBEUI Resource\ Times Exhausted

If you use the NWLink protocol, three objects are available: NWLink Internetwork Packet Exchange and NWLink NetBIOS (for computers that communicate using the IPX protocol), and NWLink Sequenced Packet Exchange (for computers that connect using the Sequenced Packet Exchange protocol). Note that frame-related counters for these objects report only zeroes.

Presentation/Program Layer

For the presentation/program layer, use the Server object counters if you monitor a server, or the Redirector object counters if you monitor a client computer. Exchange 2000, as well as some program-layer processes such as Web servers, has its own object counters that you use to monitor transmissions across this layer.

The Redirector object counters collect data about requests transmitted by the Workstation service. The Server object counters collect data about requests received and interpreted by the Server service.

At a minimum, you should include the Total Bytes Per Second counter for both the Redirector object (for client computers that you monitor) and the Server object (for server computers).

Each of these objects provides several other counters that you might want to monitor if you suspect problems with either the Workstation or Server services:

- Redirector\Current Commands
- Redirector\Network Errorslsec
- Redirector\Reads Denied/sec
- Redirector\Writes Deniedlsec
- Redirector\Server Sessions Hung
- Server\Sessions Errored Out
- Server\Work Item Shortages
- Server\Pool Paged Peak
- Server\Nonpaged Pool Failures

The Sessions Errored Out counter reports automatic disconnections and sessions that end because of an error. For more accurate values for sessions that end because of an error, obtain the value for sessions timed out and reduce the Sessions Errored Out value by that amount.

Network Diagnosis Tool

You can use the Network Diagnosis tool (Netdiag), a Windows 2000 support tool, to diagnose problems on a network. This command-line diagnostic tool helps isolate networking and connectivity problems by performing a series of tests to determine the state of your network client. These tests, and the network status information they provide, give you a more direct means of identifying and isolating network problems. Because this tool does not require you to specify parameters or switches, you can focus on analyzing the output, rather than training users how to use the tool. For more information about installing or using Netdiag, see Windows 2000 Support Tools Help.

Analyzing Performance Data

In general, performance monitoring concentrates on how the operating system and any applications or services use the resources of the system, such as the hard disks, memory, processors, and network components.

The data you accumulate through daily monitoring provides you with the information you need to analyze trends and plan your system capacity. Even if your system operates satisfactorily today, it is important to plan for changes in demand caused by new users or by technologies and programs that you deploy. Unanticipated network growth can result in overused resources and poor levels of network service. By characterizing system performance over time, you can justify the need for new resources before the need becomes critical. The following are some definitions to help you understand performance-monitoring terminology.

Throughput is a measure of the work done in a unit of time, typically evaluated from the server side in a clientlserver environment. Throughput tends to increase as the load increases, up to a peak level. It then begins to fall, and a queue might develop. Throughput in an end-to-end system, such as clientlserver, is determined by how each component performs. The slowest point in the system sets the throughput rate for the system as a whole. Often this slow point is referred to as a bottleneck. Performance monitoring identifies where bottlenecks occur in your system. The resource that shows the highest use is often the bottleneck, but not always. The bottleneck can also be a resource that successfully handles a great deal of activity. There is no bottleneck if no queues develop.

A queue is a group of jobs that are waiting to run. A queue can form under a variety of circumstances. For example, a queue can develop when requests come in for service by the resource at a faster rate than the resource's throughput, or if requests demand more time from the resource than the system can handle. A queue can also form if the requests occur at random intervals, such as large batches at the same time. When a queue becomes long, work is not handled efficiently and you might experience delays in response time.

Response ***time*** is the time required to do work from start to finish. In a client/server environment you typically measure response time on the client side. Response time generally increases as the load increases. You can measure response time by dividing the queue length for the resource by the resource throughput.

Establishing a Baseline

When you collect performance data over a period of time, with data reflecting periods of low, average, and peak usage, you can make a subjective determination of what is acceptable performance for your system. That determination is your baseline, which you can then use to detect bottlenecks or to watch for long-term changes in usage patterns that require you to increase capacity.

One of the goals of monitoring Exchange 2000 is to locate problems or anomalies in the server or network. To do this, you establish a baseline by collecting performance and diagnostic data over an extended period during varying, but typical, types of workloads and user connections. A performance baseline is a range of measurements that represents acceptable performance under typical operating conditions. This baseline provides a reference point that makes it easier to notice problems before they become serious. When you need to troubleshoot system problems, performance data gives you information about the behavior of system resources at the time the problem occurred, which is useful in discovering its cause.

When determining your baseline, it is important to know the types of work that are done and the days and times when work is done. This helps you associate work with resource usage and determine whether or not performance during those intervals is acceptable.

For example, if you find that performance diminishes somewhat for a brief period at a given time of day, and you find that many users are logging on or off at that time, it might be an acceptable slowdown. Similarly, if you find that performance is poor every evening at a certain time and you can tell that this time coincides with nightly backups when no users are logged on, performance loss might be acceptable. It is important to remember that you can make that determination only when you know the degree of performance loss and its cause.

Analyzing Results

The baseline you develop establishes the typical counter values you expect to see when your system performs satisfactorily. However, you need guidelines to help you interpret the counter values and eliminate false or misleading data that might cause you to set target values inappropriately. You need to identify and investigate bottlenecks to analyze your results and take action.

When you collect and evaluate data to establish a valid performance baseline, you should:

- Watch for unusually large values.
- Include ID threads.
- Ignore occasional spikes.
- Use graphs for reporting.
- Exclude startup events.
- Investigate zero values or missing data.

Watching for Large Values

You need to watch for values that are unusually large for one instance and not another when you are monitoring processes that have the same name. This can occur because System Monitor sometimes misrepresents data for separate instances of processes of the same name by reporting the combined values of the instances as the value of a single instance. Tracking processes by process identifier can help you solve this problem.

Including Thread Identifiers

When you are monitoring several threads and one of them stops, the data for one thread might appear to be reported for another. This is because of the way threads are numbered. For example, if you begin monitoring and have three threads—numbered 0, 1, and 2—and one of them stops, then all remaining threads are sequenced again. This means the original thread 0 no longer exists and the original thread 1 is renamed 0. As a result, data for the stopped thread 0 can be reported along with data for the running thread number 1 because old thread number 1 is now old thread number 0. To solve this problem, you can include the thread identifiers for the process in your log or display. You can use the Thread/Thread ID counter for this purpose.

Ignoring Occasional Spikes

You do not need to place too much importance on occasional spikes in data. These spikes might be due to the startup of a process and, if so, they are not an accurate reflection of counter values for that process over time. The effect of spikes can remain over time when using counters that average.

Using Graphs for Reporting

When you monitor performance over an extended period of time, you need to use graphs. Reports and histograms show only last values and averages, and they might not give an accurate picture of values.

Excluding Startup Events

Unless you specifically want to include startup events in your baseline, you must exclude them because they are temporary high values that tend to skew overall performance results.

Investigating Zero Values or Missing Data

Zero values or missing data can impede your ability to establish a meaningful baseline. You should investigate the source of these issues and obtain the missing data, if possible, before you attempt to establish a baseline.

Identifying and Investigating Potential Bottlenecks

Deviations from your baseline provide the best indicator of performance problems. However, as a secondary reference, Table 29.3 describes recommended thresholds for object counters. You can use this table to help identify when a performance problem is developing on your system.

Table 29.3 Recommended thresholds for object counters

Resource	Object\Counter	Recommended Threshold	Comments
Hard disk	LogicalDisk\ % Free Space	15 percent	None
Hard disk	LogicalDisk\ % Disk Time	90 percent	None
Hard disk	PhysicalDisk\ Disk Reads/sec, PhysicalDisk\ Disk Writes/sec	Depends on manufacturer's specification	Check the specified transfer rate for your hard disks to verify that this rate does not exceed the specifications. Some SCSI disks can handle 50 to 70 I/O operations per second.
Hard disk	PhysicalDisk\ Current Disk Queue Length	Number of spindles plus 2	This is an instantaneous counter; observe its value over several intervals. For an average over time, use PhysicalDisk\ Avg. Disk Queue Length.
Memory	Memory\ Available Bytes	Less than 4 MB	Research memory usage and add memory if needed.
Memory	Memory\ Pages/sec	20	Research paging activity.
Network	Network Segment\ % Net Utilization	Depends on type of network	You must determine the threshold based on the type of network you use. For example, for Ethernet networks, 30 percent is the recommended threshold.

Table 29.3 Recommended thresholds for object counters *(continued)*

Resource	Object\Counter	Recommended Threshold	Comments
Paging File	Paging File\ % Usage	More than 70 percent	Find the process that is using a high percentage of processor time. Upgrade to a faster processor or install an additional processor.
Processor	Processor\Interrupts/sec	Depends on processor	A dramatic increase in this counter value without a corresponding increase in system activity indicates a hardware problem. Identify the network adapter or hard disk controller card causing the interrupts. You might need to install an additional adapter or controller card. For current CPUs, use a threshold of 1,500 interrupts per second.
Server	Server\ Bytes Totallsec		If the sum of bytes total/sec for all servers is roughly equal to the maximum transfer rates of your network, you might need to segment the network.
Server	Server\ Work Item Shortages	3	If the value reaches this threshold, consider tuning the InitWorkItems or MaxWorkItems entries in the registry (in HKEY_LOCAL_MACHINE\SYSTEM\ CurrentControlSet\Services\lanmanserver\ Parameters).
Server	Server Work Queues\ Queue Length	4	If the value reaches this threshold, there might be a processor bottleneck. This is an instantaneous counter; observe its value over several intervals.
Multiple Processors	System\ Processor Queue Length	2	This is an instantaneous counter; observe its value over several intervals.

Caution Do not use a registry editor to edit the registry directly unless you have no alternative. The registry editor bypasses the standard safeguards provided by administrative tools. These safeguards prevent you from entering conflicting settings or settings that are likely to degrade performance or damage your system. Editing the registry directly can have serious, unexpected consequences that can prevent the system from starting and require that you reinstall Windows 2000. To configure or customize Windows 2000, you must use the programs in Control Panel or MMC wherever possible.

Investigating performance problems should always start with monitoring the system before looking at individual components. In precise terms, a bottleneck exists if a particular component's limitation keeps the entire system from performing quickly. Therefore, even if one or more components in your system are heavily used, if other components or the system as a whole show no adverse effects, then there is no bottleneck.

For example, suppose that a process has 10 threads, each of which uses exactly 0.999 seconds of processor time once every 10 seconds. If each thread makes a request exactly 1 second after the previous one in perfect sequence, the processor would be 99.9 percent busy. However, even though there is no queue, no interference between the threads, and, technically, no bottleneck, the system probably can't support any increased load or variation in its request scheduling without creating a bottleneck.

Factors that cause bottlenecks include the number of requests for service, the frequency at which requests occur, and the duration of each request. As long as these factors are perfectly synchronized, queues do not develop and bottlenecks do not occur. The device with the smallest throughput is typically the primary source of a bottleneck.

It is difficult to detect multiple bottlenecks in a system. You might spend several days testing and retesting to identify and eliminate a bottleneck, only to find that another bottleneck appears in its place. Only thorough and patient testing of all elements can ensure that you have found all of the problems.

It is not unusual to trace a performance problem to multiple sources. Poor response time on a workstation is most likely the result of memory and processor problems, while servers are more susceptible to hard disk and network problems.

Problems in one component might be the result, rather than the cause, of problems in another component. For example, when memory is scarce, the system moves pages of code and data between hard disks and physical memory. The memory shortage becomes evident from increased hard disk and processor use, but the problem is the lack of memory, not the processor or hard disk.

CHAPTER 30

Security

Organizations that use Microsoft Exchange 2000 Server can face many different types of security risks. Features available in Microsoft Windows 2000 Server and Exchange 2000 Server help protect your organization from many of these security risks. Windows 2000 and Exchange 2000 also provide security features to help secure your connection to the Internet. Implementing these security features is important to securing your messaging system, and can prevent disruption of messaging for your users.

In This Chapter

Security Risks

Windows 2000 Security Features

Exchange 2000 Security Features

Securing Your Internet Connection

Security Updates

Security Risks

An Exchange organization can face many different types of attacks based on a variety of different factors. These factors can range from the type of data your users send, to the type of connections your network has. Without security measures and controls in place, your data might be subject to attacks. Some attacks are passive, meaning that information is monitored; whereas others are active, meaning that the information is altered with intent to corrupt or destroy the data on the network.

Table 30.1 lists the kinds of attacks your Exchange organization might face. Implementing one or more of the security measures described in this chapter can prevent each of these attacks. You should start by identifying which of these security risks is a threat to your organization.

Table 30.1 Potential security attacks

Security Risk	Description
Identity interception	The intruder discovers the user name and password of a valid user. This can occur by a variety of methods, both social and technical.
Masquerade	An unauthorized user pretends to be a valid user. This could be a user who assumes the IP address of a trusted system and uses it to gain the access rights that are granted to the impersonated device or system.
Replay attack	The intruder records a network exchange between a user and a server and replays it at a later time to impersonate the user.
Data interception	Unauthorized users monitor and capture data moving across the network as plaintext.
Manipulation	The intruder modifies or corrupts network data. Unencrypted network financial transactions are vulnerable to manipulation. Viruses can corrupt network data.
Repudiation	Network-based business and financial transactions are compromised if the recipient of the transaction cannot be certain who sent the message.
Macro viruses	Application-specific viruses can exploit the macro language of sophisticated documents and spreadsheets.
Denial of service	The intruder floods a server with requests that consume system resources and either cause the server to stop responding or become too busy to process legitimate work. Causing the server to stop responding sometimes provides opportunities to penetrate the system.
Malicious mobile code	Malicious code runs as an auto-executed Microsoft ActiveX control or Java Applet uploaded from the Internet on a Web server.
Misuse of privileges	An administrator knowingly or mistakenly uses full privileges over the operating system to obtain private data.
Trojan horse	A malicious program that masquerades as a desirable and harmless utility.
Social engineering attack	The intruder breaks into a network by calling new employees, telling them they are from the IT department, and asking the user to verify their password for their records.

Windows 2000 Security Features

Windows 2000 Server offers you many different options for protecting your Exchange 2000 organization from the different types of possible attacks. The following are Windows 2000 Server features that relate to securing Windows 2000 Server and Exchange 2000 Server:

- Active Directory directory service
- Access Control
- Auditing
- Kerberos
- Certificate Services
- Encrypting File System
- Internet Protocol security (IPSec)
- TCP/IP filtering
- Security Configuration Tool Set

Active Directory

The Windows 2000 Server Active Directory directory service replaces the Security Accounts Manager (SAM) in Microsoft Windows NT 4.0 as a security database. On Windows NT 4.0–based servers running Microsoft Exchange Server version 5.5, there are two different sets of objects. The directory structure in Exchange 5.5 is separate from the directory structure of Windows NT 4.0. Active Directory unifies these directories to reduce administrative tasks. A security identifier (SID) essentially links Active Directory objects used in Windows 2000 and Exchange 2000 Server.

Security Subsystem

Active Directory is part of a component called the security subsystem. Discretionary access control lists (DACLs) protect all objects in Active Directory. Any attempt to gain access to an object or attribute in Active Directory is validated against the DACL by access validation routines.

Windows 2000 Server security infrastructure has four main functions:

- It is a directory service store for security policies and account information.
- It implements security models for all objects.
- It authenticates Active Directory access.
- It stores trust information for Active Directory.

Local Security Authority

The security subsystem has several components, one of which is the Local Security Authority (LSA). LSA is a protected subsystem that maintains the security for the local computer. LSA ensures that users have system access permissions. In general, LSA performs the following functions:

- It generates tokens, which contain user and group information, and the security privileges for that user.
- It manages local security policies.
- It provides interactive user authentication services.
- It manages the Audit policy and settings.

When a computer running Windows 2000 Server is configured as a member of a domain, the system's LSA uses the domain's authentication services to authenticate users. These services are in addition to the system's own local account database. The system trusts the domain to provide accurate information about security principals on the network. Figure 30.1 illustrates the components and functions of the LSA.

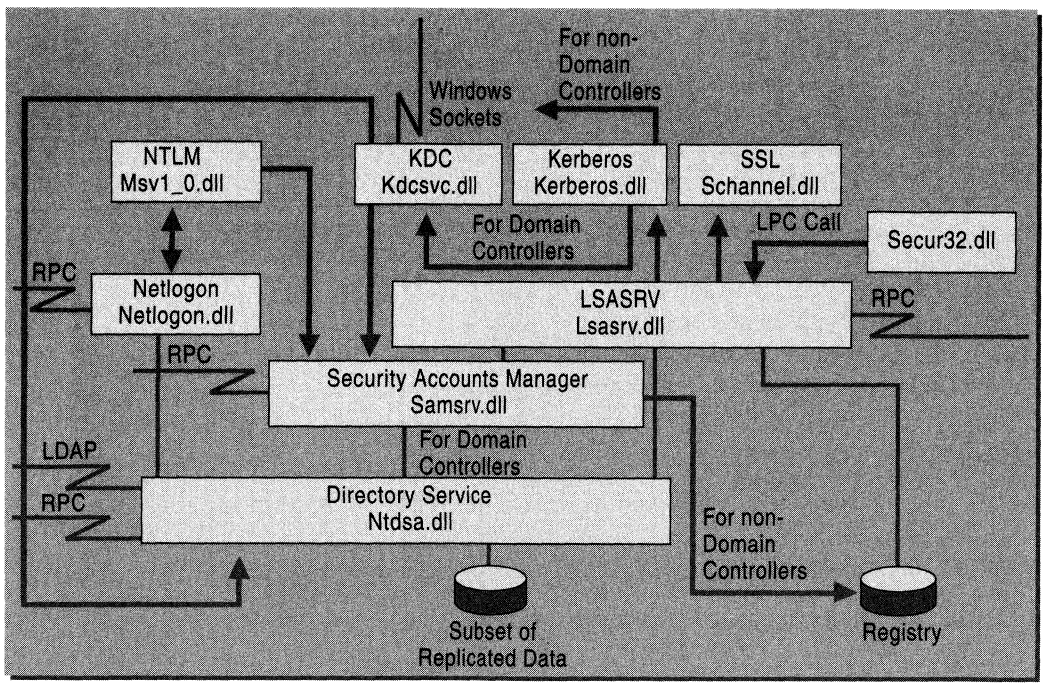


Figure 30.1 Components of the LSA

The following points describe each feature in Figure 30.1:

- **Netlogon.dll** The Net Logon service. Net Logon maintains the computer's secure channel to a domain controller. It passes the user's credentials through a secure channel to the domain controller and returns the domain SIDs and user rights for the user. In Windows 2000 Server, the Net Logon service uses Domain Name System (DNS) to resolve names to the Internet Protocol (IP) addresses of domain controllers. Net Logon is the replication protocol for Windows NT 4.0 primary domain controllers and backup domain controllers.
- **Msv1_0.dll** The NTLM authentication protocol. This protocol authenticates clients that do not use Kerberos authentication.
- **Kdcsvc.dll** The Kerberos Key Distribution Center (KDC) service, which is responsible for granting tickets to clients.
- **Schannel.dll** The Secure Sockets Layer (SSL) authentication protocol. This protocol provides authentication over an encrypted channel instead of a less-secure clear channel.
- **Kerberos.dll** The Kerberos V5 authentication protocol.
- **Lsasrv.dll** The LSA server service, which enforces security policies.
- **Samsrv.dll** The Security Accounts Manager (SAM), which stores local security accounts, enforces locally stored policies, and supports application programming interfaces (APIs).
- **Ntdsa.dll** The directory service module, which supports the Windows 2000 Server replication protocol and Lightweight Directory Access Protocol (LDAP), and manages partitions of data.
- **Secur32.dll** The multiple authentication provider that holds all of the components together.

Groups in Active Directory

Active Directory provides support for different types of groups, in addition to offering options for the scope of a group—that is, whether the group spans multiple domains or is limited to a single domain.

Two group types exist in Active Directory: security groups and distribution groups. Each of these supports one of the three group scopes: domain local, global, or universal.

The domain mode limits the choice of group type and group scope.

Group Types

The equivalent of a distribution list in an earlier version of Exchange Server is called a group in Active Directory. There are two types of groups:

- **Security groups** Groups that can be used to administer permissions for users and other domain objects. These groups are used to grant permissions on the Exchange 2000 organization to administrators and users.
- **Distribution groups** Groups that are used only for e-mail.

Group Scope

Security and distribution groups have a scope attribute. The scope of a group determines who can be a member of the group and where you can use that group in the network. Three group scopes are available: domain local, global, and universal.

- **Domain local groups** Domain local groups can have as their members groups and accounts from a Windows 2000 Server or Windows NT domain. They can be used to grant permissions only within a domain.
- **Global groups** Global groups can have as their members groups and accounts only from the domain in which the group is defined. They can be granted permissions in any domain in the forest.
- **Universal groups** Universal groups can have as their members groups and accounts from any Windows 2000 Server domain in the domain tree or forest. They can be granted permissions in any domain in the domain tree or forest.

  Note If you have multiple forests, users defined in only one forest cannot be placed into groups in another forest. Likewise, if you define groups in one forest, they cannot be assigned permissions in another forest.

Table 30.2 Security group membership rules

Domain Local Groups	Global Groups	Universal Groups
In native mode, domain local groups can have member accounts, global groups, and universal groups from any domain, and domain local groups from the same domain as members.	In native mode, global groups can have member accounts, and global groups from the same domain as members.	In native mode, universal groups can have member accounts, global groups, and universal groups from any domain as members.
Domain local groups can be put into other domain local groups and assigned permissions only in the same domain.	Global groups can be put into other groups and assigned permissions in any domain.	In native mode, universal groups can be put into other groups and assigned permissions in any domain.

Selecting a Group Scope

The global catalog maintains a list of universal group memberships. Global and domain local groups are listed in the global catalog, but their membership is not. Each change to the membership of a universal group is replicated to all global catalog servers. By minimizing the use of universal groups, you reduce the size of the global catalog, and reduce the amount of traffic on your network caused by replication of the global catalog.

- **Limit membership in universal groups to other groups** You can reduce replication network traffic by limiting membership in universal groups only to groups rather than individual accounts. To do this, you adjust the user accounts that are members of the universal group by adjusting the membership of the groups that are part of the universal group. This does not directly affect the membership of the universal group and does not generate replication network traffic.
- **Limit the use of universal groups** Limiting the use of universal groups can help you to reduce the size of access tokens when resources are in different domains. If you use global and domain local groups, the access tokens contain the global and domain local groups that are applicable to the domain in which the resource exits. If you use universal groups, the access tokens contain a list of all the universal groups to which the user belongs, even if these universal groups are not used in that domain.

Domain Local Groups

Domain local groups are best used for granting access rights to resources such as file systems or printers that are located on any computer in the domain where common access permissions are required. The advantage of using domain local groups to protect resources is that members of the domain local groups can come from both inside the same domain and outside the domain. Typically, resource servers are in domains that have trust to one or more Master User Domains, or what are known as account domains. You can use a domain local group to grant access to resources on any computer only in native mode domains. In mixed mode, domain local groups must be on domain controllers only. Table 30.3 outlines the advantages and disadvantages of domain local groups.

Table 30.3 Advantages and disadvantages of domain local groups

Advantages	Disadvantages
Membership is not published to the global catalog server, which means that changes do not incur global catalog replication.	They cannot be assigned permissions to resources in other domains.
Microsoft Outlook clients can view full user membership if they are located in the same domain as the group.	Outlook users in other domains cannot view the full memberships.
	Group membership must be retrieved on demand if expansion takes places in a remote domain.

Global Groups

You use global groups for combining users who share a common access profile based on job function or business role. Typically, organizations use global groups for all groups in which membership is expected to change frequently. The members of these groups can only be user accounts defined in the same domain as the global group. You can nest global groups to allow for overlapping access needs or to scale for very large group structures. The most convenient way to grant access to global groups is by making the global group a member of a resource group that is granted access permissions to a set of related project resources. Table 30.4 outlines the advantages and disadvantages of global groups.

Table 30.4 Advantages and disadvantages of global groups

Advantages	Disadvantages
Membership is not published to the global catalog server, which means that changes do not incur global catalog replication.	They can only contain objects from the same domain.
Outlook clients can view full user membership if they are located in the same domain as the group.	Outlook users in other domains cannot view the full memberships.
They can be used for assigning permissions to resources in the same domain.	Group membership must be retrieved on demand if expansion takes place in a remote domain.

Universal Groups

You use universal groups in larger, multi-domain organizations where you need to grant access to similar groups of accounts defined in multiple domains. It is better to use global groups as members of universal groups to reduce overall replication traffic from changes to universal group membership. You can add and remove users from the corresponding global group within their account domains. A small number of global groups are the direct members of the universal group. You can easily grant access permissions to universal groups by making them a member of a domain local group, which grants access permissions to resources. Table 30.5 outlines the advantages and disadvantages of universal groups.

Table 30.5 Advantages and disadvantages of universal groups

Advantages	Disadvantages
Membership can consist of any object in the forest.	Membership modifications incur replication to the global catalog servers.
Outlook users in any domain can view full membership.	
Membership never has to be retrieved from remote domain controllers.	

Designing a Group Implementation Strategy

The type and scope of group that you implement for Exchange 2000 depends upon your business and user requirements. To optimize flexibility, you can implement universal security groups. You can mail-enable the universal group and view it in the global address list by adding a Simple Mail Transfer Protocol (SMTP) e-mail address.

However, a disadvantage of using universal groupa is that membership is fully populated to every global catalog server, which causes replication network traffic when membership changes. For this reason, groups that might have frequent membership changes should be domain local or global groups. Because membership for these group types is not promoted to the global catalog servers, when a message is sent to the group from a remote domain, the expansion server must connect to a domain controller within the home domain for the group and retrieve the membership. This configuration has the following effects:

- The expansion server must have an IP connection to a domain controller in the remote domain.
- Retrieving membership from a remote domain can take time and slow down message delivery, which in turn can slow performance for the retrieving server.

  Note If an Exchange server already exists in the remote domain, it might be more efficient to set the expansion server to that server instead of remotely retrieving the membership. For more information about setting expansion servers, see Exchange 2000 Server Help.

You must decide whether it is better to create a universal group for the list or use domain local or global groups and accept the membership retrieval for remote domains whenever the list needs to be expanded. Consider the following for each configuration:

- Single or multiple domains in Active Directory within the organization
  - **Single** Universal groups are not necessary because all domain objects are local.
  - **Multiple** With a universal group that contains fairly static groups, users might not have access to all object attributes from other domains in universal groups.
- Whether direct IP connectivity is possible between all domains
  - **Yes** You can use universal groups for static users and include users from other domains.
  - **No** You can use dynamic lists. Users need complete access to all user attributes for users in a group. The group only contains objects from the local domain.

- Number of membership changes
  - **Frequent** If membership changes frequently, use local and global groups.
  - **Infrequently** If membership changes infrequently, use universal groups.
- Which users send the majority of mail to the list—local or remote domains
  - **From local domains** If most of the mail is sent from local domains, you should use local and global groups.
  - **From remote domains** If most of the mail is sent from remote domains, you should use universal groups.

Note When implementing domain local or global groups, Outlook users cannot view the user membership of the group unless the domain local or global groups are also contained within the local domain.

Access Control

Along with user authentication, Windows 2000 Server allows you to control access to resources, or objects, on the network. Windows 2000 Server implements access control by allowing you to assign security descriptors to objects stored in Active Directory. A security descriptor lists the users and groups that have access to an object, and the specific permissions assigned to those users and groups. A security descriptor also specifies the various access events to audit for an object.

How Access Control Works

Access control works in the following way: a program with threads of execution works on a file. The program runs as a process with threads of execution. It is actually one of those threads that opens the file. Threads are the only real agents of action on a computer.

For a thread to gain access to an object, it must identify itself to the operating system's security subsystem. A thread does not have a security identity, so it must borrow one from a security principal, such as the user. When the user logs on, the user's security identity is encapsulated in an access token that is associated with the user's logon session. When a user starts an application, the application runs as a process within the user's logon session. The application process and each of its threads of execution receive copies of the user's access token. When one of the application's threads needs to open a file, the thread identifies itself as the user's agent by presenting its access token. Thus responsibility for anything that the thread does to the file on the user's behalf is charged to the user.

Threads do not open files in the same way that users do. Threads are pieces of program code that execute on the computer, and they interact with objects through one of several APIs provided by the operating system. If a thread's job is to open a file, its code might contain the following instruction:

```
hfile=CreateFile(pszFile,GENERIC_WRITE,0,NULL,OPEN_EXISTING,0,NULL);
```

The second argument in the call to CreateFile specifies a desired set of access rights, GENERIC–WRITE, which indicates to the operating system that the thread wants to open the file and modify it. Other APIs work in a similar fashion. The caller must signal its intentions for an object by specifying the level of access that it wants.

Before allowing the thread of execution to proceed, Windows 2000 performs an access check to determine whether the security principal associated with the thread has the level of access authorization that the thread has requested. An access check compares information in the thread's access token with information in the object's security descriptor:

- The access token contains a SID that identifies the user associated with the thread and SIDs that identify the groups whose members include the user.
- The security descriptor contains a DACL with access control entries (ACEs) that specify the access rights allowed or denied to specific users or groups.

The security subsystem checks the object's DACL, looking for ACEs that apply to the user and group SIDs from the thread's access token. It steps through each ACE until it finds one that either allows or denies access to the user or one of the user's groups, or until there are no more ACEs to check. If it comes to the end of the DACL and the thread's desired access is still not explicitly allowed or denied, the security subsystem denies access to the object. Figure 30.2 illustrates this process.

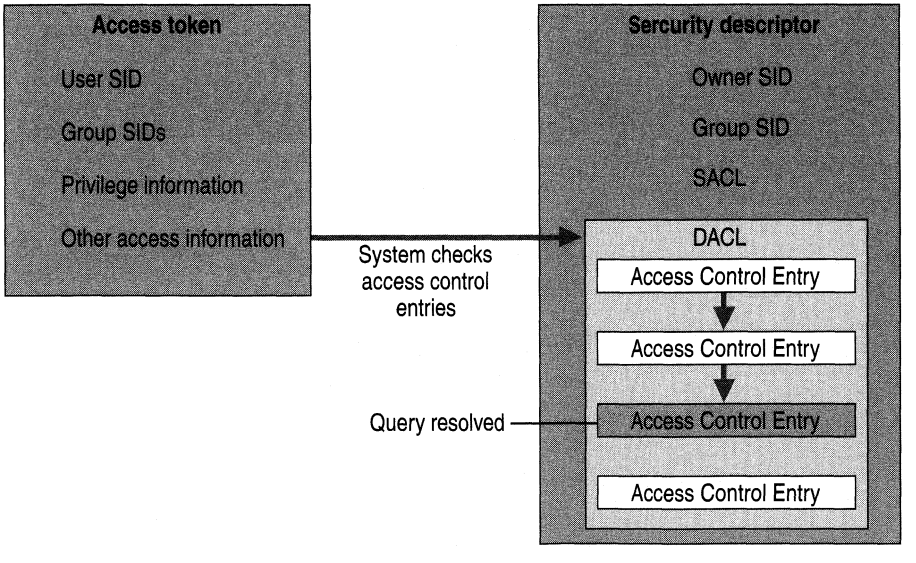


Figure 30.2 Security subsystem check

The order in which ACEs are listed in a DACL is important. For example, an object's DACL might contain one ACE that allows access to a group and another ACE that denies access to a user who is a member of the group. If access checking reaches the ACE that allows access to the user's group before it reaches the ACE that denies access to the user, the user is allowed to access the object. This is clearly not a wanted outcome.

In general, ACEs are listed in what is called *canonical order*, which places the Deny ACEs before the Allow ACEs. When the canonical order is used, an access check processes all ACEs that deny access before any ACE that allows access.

The canonical order described earlier is simplified somewhat for the purpose of this overview. It does not, for example, account for the ordering of ACEs inherited from a parent object. For the precise canonical order, see the *Windows 2000 Server Resource Kit Distributed Systems Guide*.

Exchange 2000 Access Control Model

The access control model used in Exchange 2000 Server follows that of Windows 2000 Server. The new access control model differs completely from the model used in Exchange 5.5 and allows greater access control. In Exchange 5.5, access control is based on a container-level grant, whereas Exchange 2000 allows you to set access control at the container, item, or property level.

Auditing

Security auditing is a feature of Windows 2000 Server that monitors various security-related events. Monitoring system events is necessary to detect intruders and to detect attempts to compromise data on the system. An example of an event that you can audit is a failed logon attempt.

In addition to auditing security-related events, Windows 2000 Server generates a security log and provides a way for you to view the security events reported in the log.

The Windows 2000 Server auditing feature generates an audit trail to help you keep track of all security administration events that occur on the system. For example, if you change the auditing policy so failed logon attempts are not audited, the audit trail shows this event. For more information about how to enable auditing in Windows 2000 Server, see Windows 2000 Server Help.

How Auditing Works

You can specify that an audit entry is written to the security event log whenever certain actions are performed or files are accessed. The audit entry shows the action performed, the user who performed it, and the date and time of the action. You can audit both successful and failed attempts at actions, so the audit trail can show who performed actions on the network and who tried to perform actions that are not permitted. You can view the security log in Event Viewer.

If you examine the security log regularly, you can detect some types of attacks before they succeed, such as password attacks. After a break-in, the security log can help you determine how the intruder entered the system and what they did.

Audit logging is a policy in its own right. Recording security events is a form of intrusion detection.

Auditing Exchange Active Directory Objects

Security auditing is not enabled by default. You have to activate the types of auditing you require by using the Group Policy snap-in to Microsoft Management Console (MMC).

By setting auditing on Exchange objects, such as stores and servers, you can track configuration changes to your Exchange 2000 server in the Windows 2000 Server event logs. You can audit changes in an Exchange object's configuration, permissions, or ownership. You can also perform an audit to check whether an object is written to, read, or deleted.

For more information about enabling auditing on your Exchange 2000 server, see Exchange 2000 Server Help.

Kerberos

Kerberos is the fundamental security protocol in Windows 2000 Server. Kerberos provides secure communication between Windows 2000 Server domains and clients, and is integrated into the administrative and security model.

How Kerberos Works

The Kerberos V5 authentication mechanism issues tickets for accessing network services. These tickets contain encrypted data, including an encrypted password, which confirms the user's identity to the requested service. Except for entering a password, the entire authentication process is invisible to the user.

An important service within Kerberos V5 is the Key Distribution Center (KDC). The KDC runs on each domain controller as part of Active Directory, which stores all client passwords and other account information. Figure 30.3 illustrates how the Kerberos V5 authentication process works.

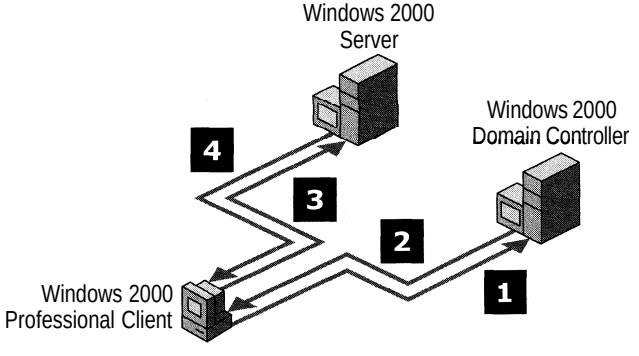


Figure 30.3 The Kerberos V5 authentication process

The Kerberos V5 authentication process performs the following steps:

1. The user on a client system, using a password, authenticates to the KDC.
2. The KDC issues a special ticket-granting ticket (TGT) to the client. The client system uses this TGT to access the ticket-granting service (TGS), which is part of the Kerberos V5 authentication mechanism on the domain controller. The TGS then issues a service ticket to the client.
3. The client presents this service ticket to the requested network service. The service ticket proves both the user's identity to the service and the service's identity to the user.
4. The user on the client successfully authenticates with the requested network service,

Kerberos and Exchange

Exchange 2000 runs as a service in Windows 2000 Server and is treated as a network service with respect to Kerberos. When the client needs to access Exchange, the client requests an Exchange service ticket from the Kerberos service. The service ticket is then used for authentication with Exchange 2000 Server. For subsequent access to the Exchange 2000 server, the client uses the service ticket, which increases authentication performance.

Kerberos Delegation of Authentication

If a client has already authenticated to a server, the server needs to make a request to another server, before the first can request the ticket on behalf of the client. The Windows 2000 Server Kerberos service supports delegation of authentication. Authentication across different domains is also possible, provided there is a trust relationship.

You can perform delegation in two ways. One way is for the client to get a session ticket for the back-end server and give it to the front-end server. Tickets obtained in this way—by a client for a proxy—are called proxy tickets. The difficulty with using proxy tickets is that the client must know the name of the back-end server. This difficulty is overcome by the second method of delegation, which allows the client to give the front-end server a TGT that the front-end server can use to request session tickets for the back-end server, as needed. Session tickets obtained in this way—with credentials forwarded by a client—are called forwarded tickets. Whether the KDC allows clients to obtain proxy tickets or forwarded TGTs is determined by the Kerberos policy set by an administrator for the domain.

Kerberos delegation is a property that you can set on the computer account in Active Directory. When the account is trusted for delegation, any computer can forward the credentials of the account. For more information about enabling a computer account for delegation, see Windows 2000 Server Help.

Delegation of Authentication in Exchange

Figure 30.4 is an example of how Exchange 2000 Server uses Kerberos delegation of authentication in a front-end/back-end Outlook Web Access environment.

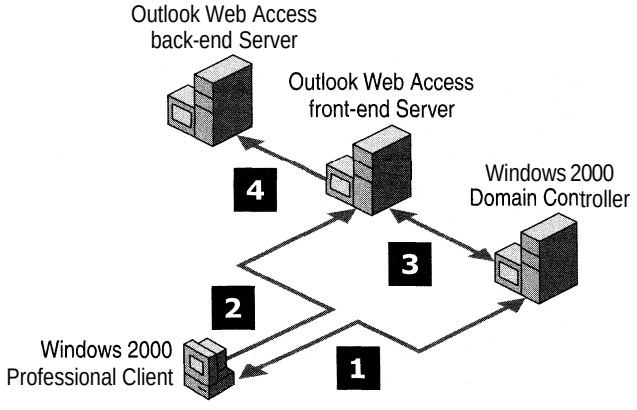


Figure 30.4 Delegation of authentication in Exchange

The Outlook Web Access deployment consists of a front-end and back-end Exchange 2000 Server, a computer running Windows 2000 Server in the same domain, and a client to make the requests. The front-end Exchange 2000 Server receives requests from clients and proxies them to the back-end server. The computer running Windows 2000 Server in the same domain functions as the Kerberos KDC. The following is an example of delegation of authentication in this environment.

1. The user on a client system, using a password, authenticates to the KDC. The KDC issues to the client a TGT that can be forwarded.
2. The client system uses this TGT that can be forwarded to access the ticket-granting service (TGS), which is part of the Kerberos V5 authentication mechanism on the domain controller.
3. The TGS then issues a forwarded service ticket to the client.
4. The user on a client system presents the forwarded ticket to the front-end server.
5. The front-end server presents the client's TGT—which can be forwarded—to the KDC. When the KDC issues a session ticket for the back-end server, it sees the FORWARDABLE flag in the TGT, sets the FORWARDED flag in the session ticket for the requested server, and returns the session ticket to the front-end server.
6. The front-end server authenticates with the back-end server on behalf of the client.

Note This example assumes that all servers are running Windows 2000 Server, and that all clients are running Windows 2000 Professional.

Certificate Services

Certificate Services provides customizable services for issuing and managing certificates used in software security systems that employ public key technologies. For more information about public key cryptography and the benefits of having a public key infrastructure (PKI), see Windows 2000 Server Help.

You can use Certificate Services in Windows 2000 Server to create a certification authority (CA), which receives certificate requests, verifies the information in the request and the identity of the requester, issues certificates, revokes certificates, and publishes a certificate revocation list (CRL). Certificate Services is a Windows 2000 Server security feature that is required to use the Key Management Service (KMS) in Exchange 2000. For more information about KMS, see "Key Management Service" later in this chapter.

Certificate Service and Active Directory

The user tabs in the Active Directory Users and Computers snap-in contain security options for each user such as Enable, Revoke, and Recover certificates.

> **Note** Every time the computer running Windows 2000 Server starts, and every 24 hours thereafter, all certificates are downloaded and added to the root store.

Encrypting File System

The Encrypting File System (EFS) is a feature of Windows 2000 Server that allows users to encrypt data directly on volumes that use NTFS file system. EFS operates by using certificates based on the X.509 standard. If no certification authority is available from which to request certificates, the EFS automatically generates its own self-signed certificates for users and default recovery agents.

The EFS supports transparent encryption and decryption of files stored on a hard disk in NTFS. A recovery policy is automatically implemented when users encrypt a file or folder for the first time. This ensures that users who lose there file encryption certificates and associated private keys can use a recovery agent to decrypt their files.

EFS is useful for mobile users who need to encrypt files or folders, such as personal mail. If the user's portable computer is lost, the data cannot be read.

Disabling EFS for All Computers in a Windows 2000 Domain

You might want to prevent users from encrypting data on their workstations. You can do this by modifying a controlling Group Policy object on domain clients, or on local computers by modifying a local Group Policy object. To disable EFS throughout a Windows 2000 domain, modify the "Default Domain Policy" Group Policy object.

To modify a Group Policy object and disable EFS throughout a Windows 2000 domain

1. Click **Start,** point to **Programs,** point to **Administrative Tools,** and then click **Active Directory Users and Computers.**
2. Right-click the appropriate node for your domain, and then click **Properties.**
3. Click the **Group Policy** tab, click the **Default Domain Policy** Group Policy object, and then click **Edit.** You do not need to use the **Default Domain Policy;** you can use a new Group Policy object such as Disable EFS to accomplish the same task.
4. In the Group Policy Editor snap-in, view the following node:

   Default Domain Policy\Computer Configuration\Windows Settings\Security Settings\Public Key Policies\Encrypted Data Recovery Agents
5. If any certificates exist in the details pane, delete them.
6. Right-click the **Encrypted Data Recovery Agents** node, click **Delete Policy,** and then click **Yes.**
7. Right-click the **Encrypted Data Recovery Agents** node, and then click **Initialize Empty Policy.**

Users on client workstations to which this policy is applied cannot encrypt files or folders. Also, if users attempt to apply encryption attributes, they receive the following error message:

```
Error Applying Attributes
An error occurred applying attributes to the file:
file name
There is no encryption recovery policy configured for this system.
```

A data recovery policy must be present to use EFS. A data recovery policy configured as Empty is not treated the same as one configured as No Policy. Deleting a policy and setting it to No Policy enables the default local policy on computers. This permits users who are local administrators to control the recovery of data on their computer. Setting a policy to Empty turns EFS off, so users cannot encrypt files on computers that fall into this category. Because policies are cumulative, enforcing an Empty policy at the domain level ensures that all Windows 2000 domain clients are denied EFS capabilities.

EFS Limitations

When an encrypted file leaves the hard disk, it is no longer encrypted. When a user sends a file as an attachment, the file is not encrypted. Files transferred over the network will require Internet Protocol security (IPSec) to maintain encryption. EFS encrypts files on the hard disk so when the directory is opened for sharing, the files cannot be read. When a legitimate user runs an application that retrieves a file from the hard disk, the file is decrypted.

Internet Protocol Security

IPSec is a framework of open standards for ensuring private, secure communications over IP networks, by using cryptographic security services. IPSec is based on standards developed by the Internet Engineering Task Force (IETF) IPSec work group.

Although Certificate Services and KMS provide security on the application layer, IPSec provides security on the IP transportation layer (that is, Layer 3). IPSec also provides protection for the TCP/IP protocol stack, such as Transmission Control Protocol (TCP), User Datagram Protocol (UDP), Internet Control Message Protocol (ICMP), and other protocols that send traffic at the IP layer. IPSec communication can transmit in blocks of data, with each block secured by a different key. This prevents an attacker from obtaining an entire communication with a single compromised key.

For complete information about IPSec and IPSec planning, see the *Windows 2000 Sewer Resource Kit TCP/IP Core Networking Guide.*

Layer 3 Protection

Security mechanisms that operate above Layer 3, such as Secure Sockets Layer (SSL), can only provide security to applications that can use SSL. An example of this type of application is a Web browser.

Using security at Layer 3 provides protection for all IP and upper layer protocols in the TCP/IP protocol stack. The primary benefit of securing information at a lower layer is that IPSec can protect all applications and services that use IP to transport data, such as Exchange 2000, without any modification to the applications or services. IPSec security at Layer 3 secures messages between Exchange 2000 servers in a way that is transparent to the servers. Figure 30.5 illustrates Layer 3 Protection.

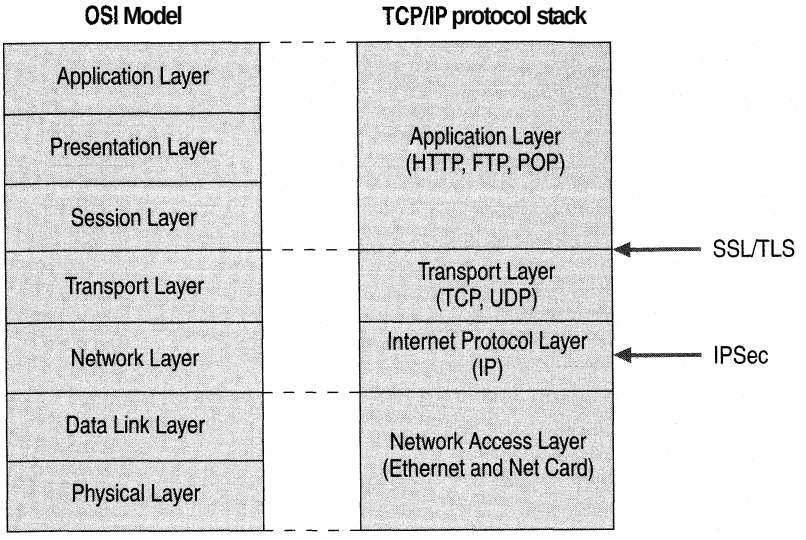


Figure 30.5 Layer 3 protection

The IPSec Model

IPSec is based on an end-to-end security model, meaning that the only computers that must know about the traffic being secured are the sending and receiving computers. Each handles security at its respective end, with the assumption that the medium over which the communication takes place is not secure. Figure 30.6 is an example Exchange organization using IPSec to secure its intranet messaging.

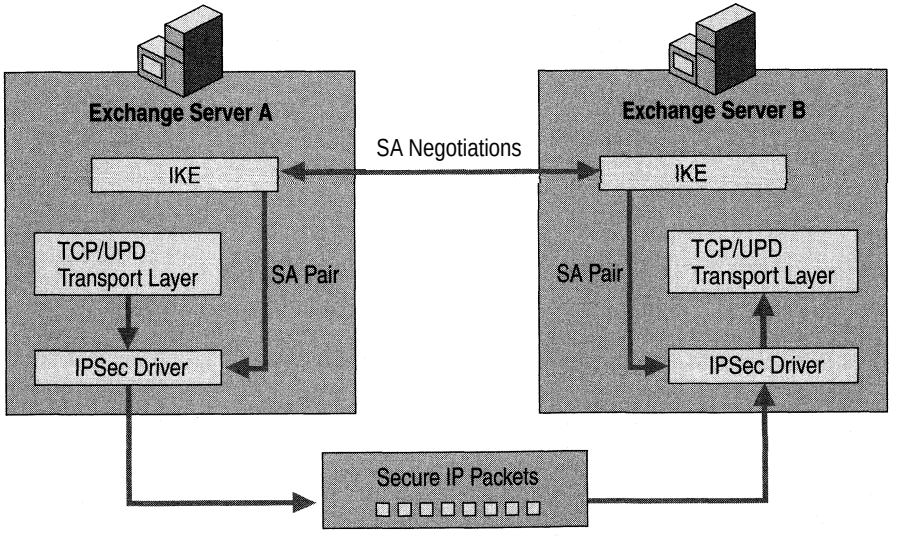


Figure 30.6 Exchange organization using IPSec to secure its intranet messaging

To secure its intranet messaging, Exchange uses IPSec in the following steps:

1. Exchange Server A sends a message to Exchange Server B.
2. The IPSec driver on server **A** checks its stored IP Filter Lists to see whether the packets should be secured.
3. The driver notifies Internet Key Exchange to begin negotiations.
4. The Internet Key Exchange service on server B receives a message requesting secure negotiation.
5. The two computers establish a Phase I security association (SA) and shared master key.
6. A pair of Phase II SAs are negotiated: one inbound SA, and one outbound SA. The SAs include the keys used to secure the information, and the Security Parameters Index.
7. The IPSec driver on server A uses the outbound SA to sign and encrypt the packets.
8. The driver passes the packets to the IP layer, which routes the packets toward server B.
9. Server B's network adapter driver receives the secure IP packets and passes them to the IPSec driver.
10. The IPSec driver on server B uses the inbound SA to check the integrity signature and decrypt the packets.
11. The driver passes the decrypted packets up to the TCP/IP driver, which passes them to the receiving Exchange 2000 Server virtual server.

Note Any routers or switches in the data path between the communicating computers will forward the secure IP packets to their destination. However, if there is a firewall, security router, or proxy server, it must have IP forwarding enabled, so that IPSec and Internet Key Exchange protocol traffic can pass through a network address translator. The Internet Key Exchange negotiation contains IP addresses in the encrypted messages, which a network address translator cannot change because the integrity hash is broken, or because the packets are encrypted.

TCP/IP Filtering

Windows 2000 Server includes support for TCPAP filtering, using the TCP/IP filtering feature. TCPIIP filtering allows you to specify exactly which types of incoming IP traffic are processed for each IP interface. This feature isolates the traffic that Internet and intranet servers process in the absence of TCP/IP filtering provided by Routing and Remote Access or other TCPIIP applications or services. TCP/IP filtering is disabled by default.

You can enable or disable TCP/IP filtering for all adapters using a single check box. Figure 30.7 shows the TCP/IP filtering dialog box, which you use to enable or disable TCP/IP filtering.

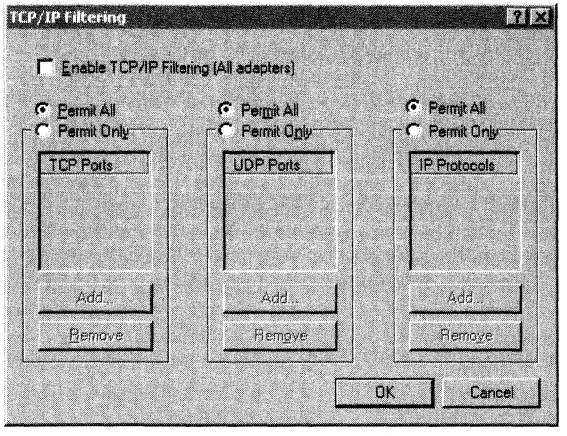


Figure 30.7 TCP/IP Filtering dialog box

You can troubleshoot connectivity problems that might relate to filtering by using the TCP/IP filtering feature. Filters that are too restrictive might not allow the kinds of connectivity that you expect. For example, if you specify a list of UDP ports and do not include UDP port 520, your computer does not receive Routing Information Protocol (RIP) announcements. This can impair the computer's ability to be a RIP router or a silent RIP host when using the RIP Listener service.

A packet is accepted for processing if it meets one of the following criteria:

- The destination TCP port matches the list of TCP ports. By default, all TCP ports are permitted.
- The destination UDP port matches the list of UDP ports. By default, all UDP ports are permitted.
- The IP protocol matches the list of IP protocols. By default, all IP protocols are permitted.
- It is an ICMP packet.

Security Configuration Tool Set

The Security Configuration Tool Set is comprised of two tools: the Security Configuration and Analysis tool, and the Security Templates tool. You can use these tools to analyze and configure security policies for your Exchange organization.

Security Configuration and Analysis uses a database to perform analysis and configuration functions. The database architecture allows you to use personal databases, import and export security templates, and combine multiple base security templates into one composite security template that you can use for analysis or configuration. The security templates feature allows you to use security templates you define or existing security templates provided with Windows 2000 Server.

The security configuration and analysis database is a computer-specific data store. The database is the starting point for security analysis and configuration. You can incrementally add new security templates to the database and create a composite security template. You can also create personal databases for storing your own customized security templates.

Security Configuration

You can use the Security Configuration and Analysis tool to directly configure local system security. The tool uses personal databases, which allows you to import security templates created with the Security Templates tool, and apply these templates to the Group Policy object for the local computer. This immediately configures the system security with the levels specified in the template.

Security Templates

Windows 2000 Server provides a centralized method of defining security with the Security Template tool. It is a single point-of-entry where you can view the full range of system security settings, which are adjusted and applied to a local computer or imported to a Group Policy object. The Security Templates tool does not introduce new security parameters—it only organizes all the existing security attributes into one place to ease security administration. You can use the Security Templates tool as a base configuration for security analysis when you use it with the Security Configuration and Analysis tool.

You can create your own security template or use a Windows 2000 predefined security template. Windows 2000 Server includes several incremental security templates. By default, these templates are stored in %SystemRoot%\Security\Templates.

You can customize these predefined templates by using the Security Templates tool and you can import them into the Security Settings extension of the Group Policy snap-in.

> **Caution** These security templates are constructed with the assumption that they apply to computers running Windows 2000 Server that use the default security settings for Windows 2000 Server. You cannot secure computers running Windows 2000 Server that are installed on file allocation table (FAT) file systems. You should not apply predefined security templates to production systems without testing to ensure that it maintains the correct level of application features for your network and system architecture.

The templates provided with Windows 2000 Server are designed to provide a starting point for configuring security based on the planned usage of the computer running Windows 2000 Server. To fully use the Security Configuration Tool Set you should do the following:

- Use the template that closest matches your organization's security policies as a base template for configuring security on Exchange 2000 servers.
- Investigate security settings specific to your organization.
- Expand your base template to include security settings specific to your organization.
- Save your updated version of the base template for quick configuration of new or restored Exchange 2000 servers.

Table 30.6 outlines the predefined Windows 2000 Server security templates.

Table 30.6 Predefined Windows 2000 Server security templates

Template	Description
Basic (basic*.inf)	The basic configuration templates can reverse the application of a different security configuration. The basic configurations apply the Windows 2000 Server default security settings to all security areas except those pertaining to user rights. The Basic template does not modify user rights settings because application setup programs commonly modify user rights to enable successful use of the application. It is not the intent of the basic configuration files to undo such modifications.
Compatible (compat*.inf)	The default Windows 2000 Server security configuration gives members of the local Users group strict security settings, while members of the local Power Users group have security settings that are compatible with Windows NT 4.0 user assignments. This default configuration enables certified Windows 2000 Server applications to run in the standard Windows 2000 Server environment for Users, while still allowing applications that are not certified for Windows 2000 Server to run successfully under the less secure Power Users configuration. However, if Windows 2000 Server users become members of the Power Users group in order to run applications not certified for Windows 2000 Server, this may be too permissive for some environments. Some organizations might prefer to assign users, by default, only as members of the Users group and then decrease the security privileges for the Users group to the level where applications not certified for Windows 2000 Server run successfully. The compatible template is designed for such organizations. By lowering the security levels on specific files, folders, and registry keys that are commonly accessed by applications, the Compatible template allows most applications to run successfully under a User context. In addition, because it is assumed that the administrator applying the Compatible template does not want these users to be Power Users, all members of the Power Users group are removed.
Secure (secure*.inf)	The Secure template implements recommended security settings for all security areas except files, folders, and registry keys. It does not modify these because file system and registry permissions are configured securely by default.

Table 30.6 Predefined Windows 2000 Server security templates *(continued)*

Template	Description
Highly secure (hisec*.inf)	The Highly secure templates define security settings for Windows 2000 Server network communications. The security areas are set to require maximum protection for network traffic and protocols used between computers running Windows 2000 Server. As a result, such computers configured with a Highly secure template can only communicate with other computers running Windows 2000 Server. They will not be able to communicate with computers running Microsoft Windows 95, Microsoft Windows 98, or Windows NT. As a result, the Highly secure templates should be used only with Exchange deployments running in native mode.
Dedicated Domain Controller (dedica*.inf)	Local user security on domain controllers running Windows 2000 Server is not secure by default. Although it is not recommended, this allows you to run existing server-based applications on domain controllers in a way that is compatible with earlier versions of Windows. If you do not run server–based applications on domain controllers, as recommended, you can define the default file system and registry permission for the local Users group in the same way as the default permissions for workstations and stand-alone servers running Windows 2000. When you implement a dedicated security template, the security settings for local Users on Windows 2000 domain controllers are applied.

Analyzing Security

The Security Configuration and Analysis tool performs security analysis by comparing the current state of system security against a security template that you import to a personal database. This template is the base configuration, and it is the template that contains your preferred or recommended security settings for that system.

Security Configuration and Analysis queries the system's security settings for all security areas in the base configuration. Values found are compared to the base configuration. If the current system settings match the base configuration settings, it is assumed that they are correct. If not, the attributes in question display as potential problems that you need to investigate.

You can create personal databases into which you can import templates for analysis. You can repeat the import process and load multiple templates. The database merges the various templates to create a composite template that resolves conflicts in the order of their importance; the most recently–imported template takes precedence when there is a conflict. After the templates are imported to the selected database, you can analyze or configure the system. For more information about the Security Configuration and Analysis tool, see Windows 2000 Server Help.

Exchange 2000 Security Features

Exchange 2000 Server provides security features specifically for securing messages. These features can secure messaging between users, help deter unsolicited e-mail, separate administrative abilities in your organization, and give users secure methods of communicating with Exchange 2000 Server. The main security features of Exchange 2000 are as follows:

- Key Management Service
- Virtual server security
- Permissions
- Securing client and server communication

Key Management Service

The KMS is a service of Exchange 2000 Server that uses Windows 2000 Certificate Services to provide secure messaging. KMS uses a variety of cryptographic technologies and methods. With the increasing need to communicate mission–critical data over public networks such as the Internet, it is important to keep data private. Cryptography is a way to protect data, thus keeping data private. However, keeping data private is only a part of cryptography. There are several other features of cryptography, which are discussed the following sections. The main KMS features are as follows:

- Encryption
- Hash functions
- Ciphers
- Algorithms
- Certificate Services and the Key Management Service

For more information about the Key Management Service, see Exchange 2000 Sewer Help.

Encryption

Encryption is the mathematical transformation of data from a readable, clear text form, into an unreadable, cipher text form. The transformation generally requires additional secret information available only to the sender and intended recipient. This information is called a ***key.*** The key allows the message to be encrypted by the sender, and decrypted only by the intended recipient. Decryption is the opposite of encryption—it transforms unreadable, cipher text data back into readable, clear text form.

Using cryptography provides not only privacy, but it provides an identity every time a user logs on to a network, accesses voice mail, or uses a user name and password to access anything. This identification is called *authentication.* Authentication is a crucial part of network data security.

As electronic transaction use increases over networks, it is important to sign documents electronically. Cryptography provides the ability to create digital signatures, which in many cases are as legally binding as a written one.

Hash Functions

A hash function provides a means of computing an electronic fingerprint, or checksum of a message. This electronic fingerprint is called the *hash* of a message.

Hashing secures messages and private key data by using them as elements in a mathematical function that creates a checksum of the package. The algorithm then is used on the receiving end to decrypt the message. Hashes typically compute quickly, and are designed so that every imaginable message can have a unique hash. Hash algorithms include MD-4, MD-5, and SHA-1.

Ciphers

A cipher is a mathematical function for encrypting and decrypting data. It is performed on readable clear text data to convert it to an unreadable version called cipher text. There are four types of ciphers: symmetric, asymmetric, block, and stream.

- **Symmetric cipher** Symmetric, or shared-key, ciphers are a form of data encryption in which a single key, known by the sender and the recipient, is used to encrypt and decrypt a message. While this form of encryption is efficient and effective, it is often difficult to share the key between both parties in a secure manner. It requires that the sender communicate the key to the recipient in a secure way.
- **Asymmetric cipher** The asymmetric cipher, or public-key cipher, is a means of solving the key management problem of symmetric key encryption. This system involves using two keys, one for encryption, and the other for decryption. One of the keys is called the public key, and the other is called the private key. You can use either the public or private key for encryption, and you use the opposite key for decryption. The public key is placed in a directory, or a location available to other users, but the private key is kept in a secure location, and is available only to the owner of the key pair. By using an asymmetric cipher, the sender and recipient do not need to agree on a key before sending data.

- **Block cipher** A block cipher uses shared-key encryption. It takes a message and breaks it into fixed length blocks, and applies the shared-key to each block. In most cases, this block size is 64 bits. The decryption operation takes the encrypted blocks, decrypts each with the same shared-key, and rebuilds the original message.
- **Stream cipher** A stream cipher is another use of symmetric encryption. Stream ciphers process small units of plaintext, usually bits. Stream ciphers are much faster than block ciphers, and can be applied to data as it is sent or received. You do not need to know the size of the message, or receive the entire message before beginning to decrypt the message. This is useful for encrypted conversations over a network such as SSL rather than individually–encrypted messages.

Algorithms

Exchange 2000 Server can use any of the following encryption algorithms:

- CAST-64
- CAST-40
- DES
- 3DES
- RC2-128
- RC2-64
- RC2-40

Certificate Services and the Key Management Service

Exchange 2000 KMS uses Certificate Services to provide security on the application layer of the messaging system. It produces X.509v3 user certificates that Exchange 2000 and Outlook 2000 use for digital signature and encryption. The X.509v3 user certificates are recognized by Secure/Multipurpose Internet Mail Extension (S/MIME) clients and ensure interoperability among different clients when used across the Internet.

> **Note** The Enroll Agent (Computer), Exchange User, and Exchange User Signature templates of the CA must be enabled before KMS installs.

KMS can be configured to require more than one administrator to be present before performing each KMS task. Each task is initially configured to require only one password, and the first password is always "password" until it is changed.

KMS implements a dual-key architecture, which gives each enrolled user two key pairs and certificates, one used for encryption, and the other used for signing. KMS uses this architecture so the signing key does not need to be archived in KMS.

Earlier versions of KMS restricted KMS to just one server per site. However, Exchange 2000 KMS fits the site-based model. The KMS organizational unit becomes the Administrator Group and each new administrator group created allows KMS to install on or point to an existing KMS server. You can also enroll users in bulk by enrolling selected users, groups, or servers.

> **Caution** When you implement encryption on the application layer, content filtering or virus checking across gateways is compromised.

Virtual Server Security

Exchange 2000 Server allows you to control which users can connect to your virtual servers by explicitly allowing or denying connections for IP addresses or domains. Exchange 2000 Server also allows you to view the transactions your virtual servers are involved with to help you identify attacks.

Virtual Server Connection Control

Controlling connections to virtual servers can prevent many of the possible network attacks listed earlier in this chapter. You can selectively include or exclude single computers, subnets, and entire domains from accessing a virtual server. By default, all computers are allowed access until you add them to a list of restricted computers. You can restrict access to a virtual server by either listing a few computers that can access the server, or by specifying a group of computers that cannot.

For example, if logging is enabled on your SMTP virtual server, you can log the IP address of the connecting client. If your Exchange server is receiving unsolicited mail from this client, for example from the IP address 172.16.0.1 as illustrated in Figure 30.8, you can deny connections to your SMTP virtual server from the client's IP address, subnet, or domain. Connection control is available on any virtual server residing on your Exchange 2000 Server except for the HTTP virtual server. For more information about how to enable connection control, see Exchange 2000 Help.

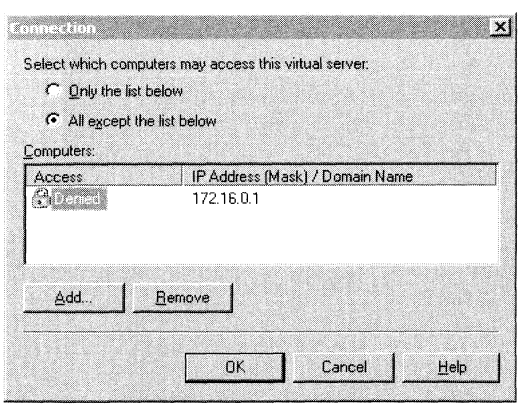


Figure 30.8 Connection Control dialog box

Protocol Logging

By setting the configuration properties of the virtual server associated with each messaging transport protocol, you can protect your e-mail system in multiple ways. The Internet protocols (SMTP, HTTP, and NNTP) enable you to use logging to track the commands the virtual server receives from clients. For example, for each message, you can see the client IP address, client domain name, date and time of the message, and number of bytes sent. Table 30.7 shows the variables that can be logged during an SMTP session.

Table 30.7 Variables that can be logged during an SMTP session

Field	Log Field	Description
Date	Date	Connection date
Time	Time	Connection time
Client IP Address	c-ip	IP address of the client that accessed the server
Client domain name	cs-username	Client that accessed the server
Service Name	s-sitename	IIS service
Server Name	s-computername	Server on which the log entry was generated
Server IP	s-ip	IP address of the server on which the log entry was generated
Method	Cs-method	SMTP protocol command sent by the client
URI Stem	cs-uri-stem	Not applicable to SMTP Service
URI Query	cs-uri-query	Varies for SMTP; depends on the SMTP protocol command
HTTP Status	sc-status	SMTP protocol reply code
Win32 Status	sc-win32-status	Windows 2000 Server status or error code; 0 indicates success
Bytes Sent	sc-bytes	Bytes sent by the server
Bytes Received	cs-bytes	Bytes received by the server
Time Take	time-taken	Length of time the action took, in milliseconds
Server Port	s-port	Not applicable to SMTP Service
User Agent	cs(User-Agent)	Not applicable to SMTP Service
Cookie	cs(Cookie)	Not applicable to SMTP Service
Referrer	cs(Referrer)	Not applicable to SMTP Service

The following is an example of log entries by an SMTP server with logging enabled. The example shows the commands used in a session to send an e-mail message. The fields used that are defined in the previous table are c-ip, cs-username, s-sitename, cs-method, sc-status, sc-bytes, and cs-bytes.

```
#Software: Microsoft Internet Information Services 5.0
#Version: 1.0
#Date: 2000-01-13 20:52:07
#Fields: c-ip cs-username s-sitename cs-method sc-status sc-bytes cs-
bytes
172.30.180.92 username SMTPSVCl HELO 250 67 12
172.30.180.92 username SMTPSVCl MAIL 250 61 49
172.30.180.92 username SMTPSVCl RCPT 250 49 47
172.30.180.92 username SMTPSVCl DATA 250 143 1268
172.30.180.92 username SMTPSVCl QUIT 0 88 4
```

When used with Windows 2000 Server event logs, the protocol log enables you to audit use of the virtual server and identify problems.

> **Note** The default Web site provided by Internet Information Services (IIS) appears in Exchange under the HTTP protocol as Exchange Virtual Server. You cannot manage this virtual server from System Manager. It must be administered in IIS.

Permissions

Exchange 2000 Server installs with various predefined permissions for groups of users that are reserved for administering Exchange and the Windows 2000 Server organizations. Within these groups are various levels of administrative permissions that can help you assign appropriate permissions to the users who are responsible for different administrative duties. Exchange 2000 Server setup also provides you alternative methods to install Exchange 2000 Server to accommodate for the typical separation of permissions between network administrators and messaging administrators found in some organizations.

Predefined Permissions

Many different types of permissions exist that you can grant on a per-user or per-group basis. Table 30.8 shows five predefined user groups with permissions already granted. These permissions have been granted on the Exchange organization.

Table 30.8 Predefined user groups with permissions already granted

Permission	Administrators	Authenticated Users	Domain Admins	Enterprise Admins	Exchange Domain Servers
Full Control	Yes	No	No	Yes	No
Read	Yes	No	Yes	Yes	Yes
Write	Yes	No	Yes	Yes	No
Create All Child Objects	Yes	No	Yes	Yes	No
Delete All Child Objects	Yes	No	No	Yes	No
Administer Information Store	Yes	No	Yes	Yes	No
Create Named Properties in the Information Store	Yes	No	Yes	Yes	No
Create Public Folder	Yes	No	Yes	Yes	No
Create Top Level Public Folder	Yes	No	Yes	Yes	No
Modify Public Folder Admin ACL	Yes	No	Yes	Yes	No
Modify Public Folder Replica List	Yes	No	Yes	Yes	No
View Information Store Status	Yes	No	Yes	Yes	No

Note Domain Admins is the Domain Admins group for the root domain in the forest. The permissions for Enterprise Admins and Domain Admins are inherited from the parent container.

Exchange Administration Delegation Wizard

The Exchange Administration Delegation Wizard helps you delegate control of Exchange configuration objects in Active Directory. By using the Exchange Administration Delegation Wizard, you can assign the task of managing different parts of Exchange to different users. You assign these tasks by assigning roles to users.

Exchange Full Administrator

The Exchange Full Administrator role grants the user permission to fully administer Exchange system information and modify permissions. Users assigned this role will have full control of the Exchange organization.

Exchange Administrator

The Exchange Administrator role grants the users permission to fully administer Exchange system information, but not modify permissions. This role may be useful for support staff who are required to administer the Exchange organization, but do not need to modify permissions.

Exchange View Only Administrator

The Exchange View Only Administrator role grants the user permission to view Exchange configuration information. This role may be useful where support staff need to view Exchange information, but do not need permission to change it.

Levels of Administration

One way to organize Administrator groups and easily grant the appropriate permissions is to create groups of administrators who have the same access privileges. The three levels of administration that should meet most organization's needs are: enterprise administrators, administrative group administrators, and recipient administrators.

Enterprise Administrators

Windows 2000 Server installs default groups in the built-in container in Active Directory Users and Computers. The built-in local security group called Administrators has all permissions to manage the Windows 2000 Server domain. The Domain Adrnins and Enterprise Admins global security groups are members of the Administrators group and therefore also are granted all permissions in the Windows 2000 domain.

The Domain Admins and Exchange Admins global security groups are granted rights to administer the Exchange 2000 organization. These rights are inherited from the parent object—the server's Configuration container.

> **Note** Exchange System Manager hides the configuration container. You can view the configuration container by running Adsiedit.exe from Windows 2000 Server Support Tools.

To assign users administrative privileges for the entire enterprise, add them to the Enterprise Admins group. By default, members of Enterprise Admins have nearly full control of both Active Directory and Exchange 2000.

Administrative Group Administrators

Many organizations might want to take advanrage of the administrative group model. To do this, you create a global security group in Active Directory and grant this group one of the roles in the Exchange Administration Delegation Wizard for the specific administrative group. These permissions should be the same as those for Enterprise Admins, except that they are only valid within the selected administrative group,

Recipient Administrators

Recipient Administrators administer all aspects of user objects. You can use the built-in Windows 2000 Server Account Operators security group as a single location for recipient administrators. You should grant the Account Operators group Exchange View Only permissions role using the Exchange Administration Delegation Wizard. Recipient administrators must be able to create accounts in Active Directory in addition to enabling a mailbox in Exchange 2000.

All user administration permissions must include rights to Active Directory in addition to Exchange. This reflects a change from earlier versions of Exchange where Exchange managed its own directory rather than relying on the operating system,

> **Note** Any user that you want to administer any level of Exchange 2000 must have at least Read permissions on the Exchange organization container.

Organlratlon Preparation

Preparing your organization for Exchange 2000 involves two setup options: ForestPrep and DomainPrep, You can prepare Windows 2000 Server for the setup of Exchange 2000 Server by separating parts of Exchange 2000 Setup that require high level network access permissions from those that don't.

Typically, organizations with messaging systems have two sets of administrators: network administrators and messaging administrators. These two sets of administrators typically do not share the same permissions. To allow for this separation, Exchange 2000 Setup provides ForestPrep, and DomainPrep.

The following are reasons to use ForestPrep and DomainPrep:

- If you want a user that doesn't have high level network permissions to install the first Exchange 2000 Server.
- If you want to resolve replication latency issues, and speed up the installation of Exchange 2000 Server.

Using ForestPrep and DomainPrep is not necessary when all of the following conditions are true:

- If all Exchange 2000 Servers are installed in a single domain.
- If the single domain contains the schema master.
- If all Exchange users reside in that single domain.
- If the account installing Exchange 2000 Server has Enterprise and Schema Administrator permissions.

ForestPrep

The ForestPrep setup option performs all the tasks that require Enterprise and Schema Administrator permissions. It is designed to be used by the highest network administrator. and is only required to run once per forest.

During ForestPrep setup, you can designate the Exchange 2000 Administrator account. When ForestPrep and DomainPrep are complete, the permissions on this account are equivalent to the Delegation Wizard Org-Level Exchange Full Administrator account.

Before running **setup /ForestPrep** you should do the following:

- Establish the Exchange 2000 Administrator account.
- Determine if you want to create a new Exchange 2000 organization or join an existing Exchange *5.5* Server organization.

  Caution This decision is permanent and cannot be reversed.
- Establish a name for the organization if you are not joining an existing organization.
- If you are joining an existing Exchange *5.5* organization, you must know the name of the Exchange *5.5* server and the service account and password in the site you want to join.

The ForestPrep requirements are as follows:

- The ForestPrep setup option must be run in the same domain as the Schema Master.
- The user running the ForestPrep setup option must be granted Enterprise and Schema Administrator permissions.
- If the user running the ForestPrep setup option is joining an existing Exchange *5.5* organization, this user must also have administrator permissions on the Exchange *5.5* Site container and Configuration container.
- If joining an existing Exchange *5.5* organization, Active Directory Connector and Exchange Server version 5.5 Service Pack 3 (SP3) must already be installed.

The following describes a ForestPrep installation:

- It extends the Active Directory schema.
- It creates the Organization container and the Global containers underneath it.
- If joining an existing Exchange 5.5 organization, ForestPrep also replicates objects from the directory service of the Exchange 5.5 server to which you chose to bind.
- It gives the account specified the same permissions on the MSExchange container as a Delegation Wizard Org-Level Exchange Full Administrator.

Note After ForestPrep completes. you must still run DomainPrep.

Table 30.9 illustrates when running ForestPrep is required, and when it is not required.

Table 30.9 When ForestPrep is required

Running ForestPrep is Required	Running ForestPrep is Not Required
When installing the first Exchange 2000 server into the Windows 2000 Server organization into a domain that does not contain the Schema Master.	When installing the first Exchange 2000 Server into the Windows 2000 Server organization into a domain that contains the Schema Master.
When an account that doesn't have Enterprise, Schema, or Domain Administrator permissions wants to install the first Exchange 2000 server.	When an account that has Enterprise, Schema, or Domain Admins permissions wants to install the first Exchange 2000 server. This account will automatically become the Exchange 2000 Administrator account.

DomainPrep

The DomainPrep setup option performs the setup tasks that require Domain Administrator permissions. It is designed so that only users with Domain Administrator permissions can use it. You should run **setup /DomainPrep** in each domain that will have an Exchange 2000 server installed in it. You should also run it in each domain that contains users with Exchange mailboxes. The user running DomainPrep is not required to be an Exchange administrator.

The following are DomainPrep requirements:

- The user running the DomainPrep setup option must have Domain Administrator permissions for the domain.
- The ForestPrep setup option must run before DomainPrep
- Replication from ForestPrep must complete.

Note Exchange Administrator and Exchange 5.5 permissions are not required to run DomainPrep.

The following lists what the DomainPrep installation accomplishes:

- It creates two new domain groups: Exchange Domain Servers and Exchange Enterprise Servers.
- It creates the PF Proxy container.
- It grants permissions to the Exchange 2000 Administrators and Exchange Servers on these objects.

Table 30.10 illustrates the differences between the Exchange Domain Servers group and the Exchange Enterprise Servers Group.

Table 30.10 Differences between Exchange Domain Servers group and Exchange Enterprise Servers group

Exchange Domain Servers Group	Exchange Enterprise Servers Group
Glabal Security Group	Domain Local Security Group
Contains the computer accounts for all the Exchange servers in the domain. This group is populated by Setup during the regular installation process, not during DomainPrep.	Contains the Exchange Domain Servers groups from all the domains running Exchange 2000 Server. DomainPrep adds the Exchange Domain Servers group from the current domain; the Recipient Update Service adds the domains that have an active Recipient Update Service.
The Exchange Domain Servers group is needed for the Recipient Update Service.	Exchange Enterprise Servers group is granted permissions on the domain container.

Table 30.11 illustrates when running DomainPrep is required and when it is not required.

Table 30.11 When DomainPrep Is required

Running DomainPrep Is Required	Running DomainPrep Is Not Required
After running ForestPrep.	When the account installing Exchange 2000 Server is granted Exchange Full Administrator permissions and Domain Administrator permissions. In this case, DomainPrep runs automatically.
Before using a designated Exchange 2000 administrator to install the first Exchange 2000 Server into a domain.	
When you want to create a Recipient Update Service for a domain that isn't going to host an Exchange 2000 Server, but will have mail-enabled users.	

Active Directory connector

Active Directory Connector requires the administrator to be granted specific permissions when performing actions such as installing the connector and creating connection agreements. The following are the required permissions for performing specific Active Directory Connector administrative functions.

Active Directory Connector Installation

Installing the first Active Directory Connector in a Windows 2000 Server forest involves extending the Active Directory Schema in addition to copying files to the local computer, Therefore, the administrator performing the installation must be a member of the following groups:

- Schema Admins
- Enterprise Admins
- Administrators (of target computer)

After the first Active Directory Connector installs, the Active Directory Schema is updated. This eliminates the need for the user to be a member of the Schema Admins group. The user performing the installation must be a member of the following groups:

- Domain Admins (local domain)
- Administrators (of target computer)

Active Directory Connector Service Account

When you install Active Directory Connector, you must have a service account. This is because some of the technology required for the Active Directory Connector is part of Windows 2000 Server. Exchange 2000 Server has features to prepare Active Directory for installation of the server. Part of this preparation involves setting permissions for LocalSystem services to Active Directory. Because you can use the Active Directory Connector without Exchange 2000 Server installed, a separate service account is used.

The Active Directory Connector service account requires the following permissions:

- Member of the Built-in\Administrators group
- Member of Enterprise Admins if used only with Windows 2000
- Member of Enterprise Admins or Exchange Full Administrator role if used with Exchange 2000 Server

Connection Agreement Credentials

When you create a connection agreement in Active Directory Connector, credentials for accessing Active Directory and the Exchange directory are required. The account you provide should have the following permissions:

Exchange 5.5

- Exchange Administrator role to the Exchange *5.5* Site (local) naming context
- Exchange View Only Administrator role to the Exchange *5.5* Site (remote) naming context

  Exchange View Only Administrator role to the Exchange *5.5* Organization naming context

Active Directory

Domain Admins (local domain)

Exchange View Only Administrator role to the Exchange 2000 organization

Other Permissions Issues

Like Active Directory Connector, there are specific Exchange 2000 administrative functions that require specific permissions to be successfully performed. The following are Exchange functions and administrative actions that require various permissions to perform successfully.

Creating and Deleting Mailboxes

To create a mailbox, you need to have permission to create a user in Active Directory. For example, you can be a Domain Administrator, or you can have delegated access to a specific organizational unit. If you have delegated access. you must have the Exchange View Only Administrator role to the Administrative Group where the target Exchange 2000 Server exists.

Upgrading Mailboxes In The Same Site/Administrative Group

To move a mailbox from Exchange *5.5* to Exchange 2000 using the Active Directory Users and Computers snap-in, you must have the following permissions:

- **Exchange 5.5** Administrator privileges to the Site naming context.
- **Active Directory** Domain Admins or Server Operators group permissions in the local domain.

Configuring Routing Groups and Connectors

Because configuring routing groups and connectors does not directly affect user accounts, you only need Exchange Administrator permissions for the Administrative Group where the target routing group exists.

If you need to define the global message formats for specific outbound domains or need to specify global message thresholds, you need Exchange Administrator permissions for the Exchange Organization.

Manipulating Message Queues

To view message queues in Exchange System Manager, you need the following permissions:

- Exchange View Only Administrator on the administrative group where the connector exists.
- Member of the local Administrators group on the target computers.

To remove messages from queues, you need the following permissions:

Exchange Administrator role on the administrative group where the connector exists.

- Member of the local Administrator group on the target computers.

Installing Exchange 2000 Server with Exchange Administrator Permissions

The following are issues you should consider when installing Exchange 2000 Server with Exchange permissions:

- The Exchange Administrator must first be an Administrator of the local computer on which Exchange 2000 Server is installed.
- If joining an Exchange 5.5 organization, the Exchange Administrator must also have administrative permissions on the Exchange 5.5 Site and Configuration containers.
- The Exchange Administrator must also know the Exchange 5.5 service account and password.

Exchange 2000 Administrator Limitations

There are certain actions that a user granted Exchange Administrator permissions is not allowed to perform:

- Users granted only Exchange Administrator permissions cannot run ForestPrep or DomainPrep.
- The Exchange Administrator cannot create new users or give users Exchange mailboxes unless the user also has Account Operators permissions.

Securing Client and Server Communication

Exchange 2000 Server allows you to secure communication between messaging clients and your servers. Securing communication between the client and server can prevent interception of sensitive e-mail, and help prevent users from sending and receiving unsolicited mail.

Inbound Encryption

You can configure the Exchange 2000 protocol virtual servers such as SMTP servers and Post Office Protocol version 3 (POP3) servers to require inbound encryption when a client is communicating with the server. You can configure each server to use SSL with basic authentication, further securing your messaging environment. Requiring security on your virtual

server, especially your SMTP virtual server, makes it more difficult for users to send and receive unsolicited e-mail. For more information about inbound encryption, see Exchange 2000 Help for the protocol virtual server you want to secure.

Encrypted RPC

You should configure your client to use encrypted remote procedure calls (RPCs). This ensures that messages transmitted over the Internet between clients and servers are secure and no users can tamper with them. Exchange 2000 Server uses RPC when communicating between the client and the server. Exchange uses the security built into RPC to authenticate client-server and server-server communications.

Although Exchange 2000 Server uses SMTP as its native transport protocol between Exchange 2000 servers, there are situations when it uses RPC. You can encrypt RPC to protect your client-server communication in the following circumstances:

- If a MAPI client with encryption enabled connects to the Exchange 2000 server.
- If an Exchange 5.5 server connects to another server in the same site or over Site Connector.
- If an Exchange 2000 server connects to an Exchange 5.5 server, whether it is in the same mixed site and routing group, or over Site Connector and Routing Group Connector.

You can encrypt the entire client-server communication for secure mailbox access over the Internet.

Encrypted RPC uses a 40-bit RSA algorithm called RC4 to encrypt data while it is on the network. You can configure Outlook to use encrypted RPC so communication between clients and servers is secure and no users can tamper with messages during transit.

Encrypting RPCs is different from encrypting a message using advanced security encryption; it provides protection for data only while it travels from point to point on the network. A message encrypted using advanced security encryption is protected until the recipient decrypts it using the client, regardless of how many hops are used during delivery. Encrypted RPCs provide increased security for messages sent on internal networks, as well as to outside organizations on the Internet.

To configure encrypted RPCs

1. In Microsoft Outlook, on the **Tools** menu, click **Services.**
2. In the list of information services, click **Microsoft Exchange Server,** and then click **Properties.**
3. Click the **Advanced** tab.
4. Under **Encrypt information,** select both check boxes to encrypt all client/server communication.

Securing Your Internet Connection

Internet connections expose your Exchange 2000 servers to traffic on the Internet, and increase the possibility of the security risks listed earlier in this chapter. You should use virus security and physical security to secure your Internet connection.

Virus Protection

Viruses can enter a computer system through the Internet, newsgroups, or e-mail. Once in the system, they can spread when e-mail, files, or documents are shared. There are two ways to protect against viruses. One is to install virus-scanning software, which can detect and remove viruses, on the computers in your organization. However, you must continually update your virus-scanning software to protect your computers as viruses evolve. The other way to protect against viruses is to prevent the virus from entering your system by using firewall software that includes virus detection. You must continually update your firewall software to keep up with changing viruses.

Even if you have virus-detection software in the firewall and on desktop computers, viruses can still enter the system. Messages that are encrypted cannot be scanned for viruses. A virus contained in an encrypted message escapes detection and can be activated when a user in the organization opens the message. For maximum protection, you should ask users to minimize their exposure to viruses by adhering to the following practices:

- Tell users to be cautious with attachments, and never to open attachments from unknown sources.
- Tell users to log off of their e-mail accounts when not using them. If they log on remotely from a public terminal, tell them to log off when they end their session.
- Tell users that when they respond to unsolicited e-mail it only confirms that they have an active e-mail address and it can expose your organization to attack. Ask users to forward unsolicited e-mail to you so you can filter it out by blocking mail from those IP addresses.
- Tell users not to respond to requests for personal information, such as their passwords, even if the request seems to come from someone reputable.

Physical Security

Aside from using the security features provided in Exchange 2000 and Windows 2000 Server, you can implement a variety of physical security measures to protect your Exchange organization.

Firewalls

Firewalls are one of the best ways to protect your systems from attacks by users on the Internet. You can use a firewall to separate your internal network from the Internet. A firewall restricts inbound and outbound access, and it can analyze all traffic between your network and the Internet. A firewall can range from a simple packet filter to complex bastion hosts that analyze traffic for each application type. A bastion host must be secured because it is accessible from the Internet and exposed to attack. A firewall can be a single router or computer, or it can be a combination of components such as routers, computers, networks, and software such as Microsoft Proxy Server.

A firewall or proxy server fundamentally separates one network from the other, such as the connection of the LAN to the Internet. For a list of ports and protocols used by Exchange and Windows 2000 Server, see Appendix B, "Ports and Protocols." It is important to review these ports and protocols to prevent interruption of service when implementing a firewall.

Dual-Homed System

One way to set up a bastion host is to use a dual-homed computer, which has a connection to two networks but does not route packets between them. One of the connections is to your internal network and allows communication with other servers and clients in your organization. The other connection is to the Internet. You can run Exchange on a dual-homed computer to provide safe e-mail connectivity to the Internet.

Proxy Sewers

Some services, such as Web and File Transfer Protocol (FTP), are point-to-point, so a client can make a connection directly to a server. Allowing clients inside your network to connect directly to hosts on the Internet is generally unsafe. One solution to this problem is to use a proxy server to interact with external servers on the client's behalf. The client communicates with the proxy server, which relays approved client requests to servers and relays responses back to the client. External hosts do not connect directly to clients in your network.

Domain Name System

If your system accepts mail directly from other hosts on the Internet, is should be listed in the Domain Name System (DNS). When you register your system with DNS, a DNS Mail Exchanger (MX) record is created that routes all mail to your host that processes incoming mail for the domain. Unless you plan to forward all outbound Internet mail to a relay host (a host outside your organization that has better e-mail connectivity), your server must be able to query DNS to deliver messages. You can configure your Exchange Server 2000 server to use DNS services from your Internet service provider (ISP), or you can use your own DNS servers. If you maintain your own DNS servers, they must be registered with your parent domain.

If you are using DNS and do not want DNS queries from the Internet to return information about computers on your internal network, configure DNS so that external hosts can query information about your internal servers but not about other hosts. To do this, your must set up a pair of DNS servers: an external DNS server for your bastion hosts, and an internal DNS server that clients on your network use. Configure the internal DNS server to forward queries it cannot resolve to the external DNS server so clients in your network can resolve Internet host names. Your bastion host should use the internal server for DNS to resolve both internal and external names. Because the external DNS server does not have complete information for your internal network, and because access to your internal DNS server is not available from the Internet, you can hide most of your computers from external DNS queries by not creating records for them on the external DNS server

Security Updates

To keep up with updates and security notices released by Microsoft:

- Make sure that you have the latest service packs and hot fixes. You can download the latest updates from http://www.microsoft.com. For more information about Windows updates, see Windows 2000 Server Help.
- Subscribe to the Microsoft Security Bulletin at http://www.microsoft.com.

CHAPTER 31

Optimizing Exchange 2000

Optimizing your Microsoft Exchange 2000 Server organization is the end result of planning, trial and error. and specific steps. An optimized Exchange 2000 organization is one that provides the best features to users with the lease administrative effort and costs.

Optimizing an Exchange 2000 organization involves optimizing the Active Directory directory service, which includes customizing your forest and establishing a logical structure for your Exchange organization. The Active Directory schema is the directory in which users search for and interact with other users. You should tune your Exchange 2000 servers to provide the best connectivity to your users. There are many ways to plan and create your Exchange message transport topology for optimal performance.

In This Chapter

Optimizing Active Directory

Active Directory Schema

Optimizing Exchange 2000 Server

Tuning Exchange 2000 Performance

Optimizing Message Transport

Optimizing Active Directory

Active Directory is the central repository of information in your Exchange 2000 organization. Properly optimizing Active Directory helps improve performance, ease administration, and keep your organization running smoothly.

Optimizing Active Directory includes customizing forests by using trees and adding domains, using organizational units to partition administrative privileges, establishing trust relationships and domain controller roles, and creating sites based upon your underlying network infrastructure.

The following sections present overview information about Active Directory, and how you can optimize it.

- Forests
- Trees
- Domains
- Trust relationships
- Domain controller roles
- Sites

For more detailed information about Active Directory, see Microsoft Windows 2000 Server Help.

Forests

Forests contain trees and domains, and are the boundary for Active Directory. You can customize forests to meet your business needs, and you can administer them in many different ways.

You optimize your forest by using either a single forest or multiple-forest environment. For more information about forests, see Windows 2000 Server Help.

Single-Forest Environment

A single forest environment is simple to create and maintain. All users see a single directory through the global catalog, and do not need to be aware of any directory structure. When adding a new domain to the forest, no additional trust configuration is required. You only need to apply configuration changes once to affect all domains.

> **Note** You can only deploy Exchange 2000 over a single forest. An Exchange 2000 organization cannot span multiple forests. You can deploy many trees in a single forest, each having access to the Exchange 2000 organization. If more than one forest is required, you must have a separate deployment of Exchange 2000 Server for each forest.

Multiple-Forest Environment

If administration of your network is distributed among many autonomous divisions, it might be necessary to create more than one forest. If you choose to create multiple forests, you should understand the implications of having multiple forests and the limitations it imposes on users and objects.

Because forests have shared elements, such as schemas, all the participants in a forest must agree on the content and administration of those shared elements.

Limitations

Exchange 2000 Server requires Active Directory. If you need to deploy two domains in two separate forests, a single Exchange 2000 organization cannot span this deployment because each forest maintains its own Active Directory.

The dependency of Exchange 2000 Server on Active Directory yields the following results in this scenario:

- You must administer two separate Exchange organizations.
- There is no automatic directory replication between the two organizations, so you have two separate global address lists.
- You cannot use Routing Group connectors between the two organizations. You must use Simple Mail Transfer Protocol (SMTP) connectors or X.400 connectors.
- There is no link state data transfer for Exchange 2000 because you cannot use Routing Group connectors.

Note When two Microsoft Windows NT 4.0 domains have no trusts between each other, you can deploy a single Microsoft Exchange Server 5.5 organization across the two domains because there is no reliance on the underlying security structure. Because Exchange Server 5.5 performs e–mail–based directory replication between sites, users in both domains can see all users within the same organization by using a global address list.

Trees

Trees are a hierarchical arrangement of Windows 2000 Server domains that share a contiguous namespace. You form trees in your Windows 2000 Server forest when separate groups of servers require separate administration. The parent-child relationship is a naming and trust relationship only. Administrators in a parent domain are not administrators of its child domain.

Forming trees can separate administration of Exchange 2000 servers. For example, if you have a group of servers in your organization that require a separate group of administrators, forming a tree in the forest made of these servers allows a separate group of administrators to maintain the servers. By using the Exchange System Manager snap-in, you can create a new administrative group and routing groups to manage this new tree.

Domains

The basic unit of logical structure in Active Directory is the domain. A domain is a collection of computers defined by an administrator that share a common directory database.

To optimize your domains, determine the number of domains you require in your tree or forest, decide whether to preserve or upgrade existing Windows NT 4.0 domains (if you use Windows NT 4.0), and partition administrative privileges by using organizational units.

DeterminIng the Number of Domains

To determine the number of domains that you want in each forest, start by considering a single domain only, even if you currently have more than one domain. Next, provide a detailed justification for each additional domain. Every domain that you create introduces incremental cost in terms of additional management time. For this reason, be certain that the domains you add to a forest serve a beneficial purpose.

> **Note** If you already have domains running Windows NT, you might want to leave them as they are instead of consolidating them into a smaller number of domains in Active Directory. Whether you decide to keep or consolidate a domain, be sure to weigh those costs against the long-term benefits of having fewer domains.

Some reasons to create more than one domain include:

- Different password requiremenrs between organizations
- Large numbers of objects
- Different Internet domain names
- Better control of replication
- Decentralized network administration
- Preserving existing Windows NT domains
- Administrative partitioning
- Physical partitioning

Administrative Partitioning

More domains might be necessary depending on the administrative and policy requirements of your organization.

If there are unique domain user security policy requirements, you might want to have a set of users on your network abide by a domain user security policy that is different from the security policy applied to the rest of the user community. For example, you might want your administrators to have a stronger password policy, such as a password change interval that is shorter than the regular users on your network. To do this, you must place those users in a separate domain.

If a division requires autonomous domain administration supervision, the members of the domain's Administrators Group in that domain have complete control over all objects in that domain. If you have a division in your organization that will not allow outside administrators to have control over their objects, place those objects in a separate domain. For example, for legal reasons, it might not be advisable for a subdivision of an organization that works on highly sensitive projects to accept domain supervisions from a higher-level group. Remember that all domains in the forest must share the configuration container and schema.

Physical Partitioning

Physical partitioning involves dividing the domains you have in a forest into a number of smaller domains. Having a greater number of smaller domains allows you to optimize replication traffic by replicating objects only to places where they are most relevant. For example, in a forest containing a single domain, every object in the forest is replicated to every domain controller in the forest. This might lead to objects being replicated to places where they are rarely used, which is an inefficient use of bandwidth. For example, a user that always logs on at a headquarters location does not need the user account replicated to a branch office location. Creating a separate domain for the headquarters location and not replicating that domain to the branch office can avoid replication traffic.

Organizational Units

An organizational unit is a container object that you can use to organize objects within a domain. An organizational unit contains objects, such as user accounts, groups, computers, printers, and other organizational units.

To optimize your domains by using organizational units, you can group objects into a logical hierarchy based on your organizational structure, delegate administrative control over organizational units to partition administrative privileges, and simplify an existing multiple domain model into a single domain model.

Hierarchy

You can use organizational units to group objects into a logical hierarchy to represent an organization's organizational structure based on departmental or geographical boundaries. You can also use them to create a network administrative model based on administrative responsibilities. For example, you can create two organization units, one for an administrator who is responsible for all of the user accounts and another for the administrator who is responsible for all of the computers. To do this, you create one organizational unit for users and another for computers.

Each domain can implement its own organizational unit hierarchy. The organizational unit hierarchy within a domain is independent of the organizational unit hierarchy of other domains.

Administrative Control of Organizational Units

You can delegate control over the objects within an organizational unit. To delegate administrative control of an organizational unit, you grant specific permissions for the organizational unit and the objects that it contains to one or more users and groups.

For an organizational unit, you can assign full administrative control over all objects in the organizational unit, or you can assign limited administrative control, such as the ability to modify e-mail information on user objects in the organizational unit.

Organizational Units and the Single Domain Model

Because an Active Directory domain can contain millions of objects, you can convert from a multiple domain model to a single domain model. This simplifies management at the domain level. You can combine domain resources into organizational units that best meet your organization's requirements, rather than creating and administering multiple domains. You can easily move objects between organizational units within the domain, nest organizational units within each other, and create new organizational units as the need arises.

Trust Relationships

Active Directory supports two forms of trust relationships: one-way non-transitive trusts and two-way transitive trusts. Windows 2000 Server optimizes trusts between domains by using two-way transitive trusts.

One-Way Non-Transitive Trusts

In a one-way trust relationship, if domain A trusts domain B, domain B does not trust domain A. In a one-way non-transitive trust relationship, if domain A trusts domain B, and domain B trusts domain C, domain A does not trust domain C. Figure 31.1 illustrates this type of trust relationship:

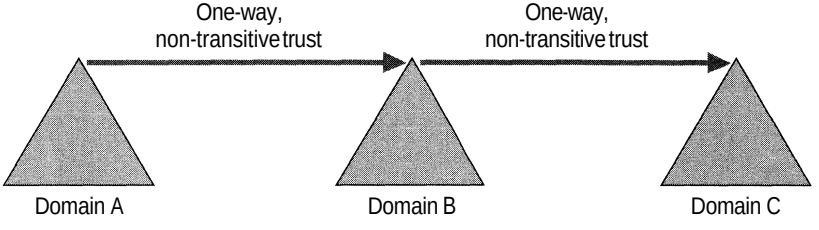


Figure 31.1 One-way non-transitive trust

> **Note** Windows NT networks use one-way non-transitive trust relationships. You manually create one-way non-transitive trust relationships between existing domains. In a large network, this form of trust relationship creates a large amount of administrative work.

Active Directory supports one-way non-transitive trusts for connections to Windows NT networks. You can also establish one-way non-transitive trusts between Active Directory domains. For example, if you want to allow an external business partner to have access to resources in a particular domain while working on a joint project, you might create a one-way non-transitive trust between the internal and external domains.

Two-Way Transitive Trusts

In a two-way trust relationship, if domain A trusts domain B, domain B trusts domain A. In a two-way transitive trust relationship, if domain A trusts domain B and domain B trusts domain C, domain C trusts domain A. Figure 31.2 illustrates this type of trust relationship:

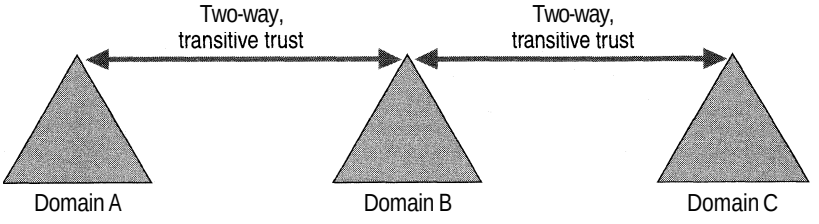


Figure 31.2 Two-way transitive trust

If a two-way transitive trust exists between two domains, you can grant permissions to resources in one domain to user and group accounts in the other domain, and vice versa. Two-way transitive trust relationships are the default relationship between Windows 2000 domains.

Limiting Trust Relationships

Because all domains in a forest trust one another by default, every user in the forest can be included in a group membership or appear on an access control list (ACL) on any computer in the forest. If you want to prevent certain users from receiving permission to certain resources, those users must reside in a different forest from the resources. If necessary, you can use explicit trusts to allow those users to receive access to resources in specific domains.

Domain Controller Roles

You can assign special roles, such as global catalog server, and operations to specific domain controllers. It is important to understand these roles because if one of the servers that serves that role is not available, the role's function is not available in Active Directory. The domain controller roles are as follows:

- Global catalog
- Operations master
- Schema master
- Domain naming master
- Relative identifier (RID) operations master
- Primary domain controller (PDC) emulator
- Infrastructure master

Global Catalog

The global catalog is a repository of information that contains a subset of attributes for all objects in Active Directory. By default, the attributes that are stored in the global catalog are those that are most frequently used in queries (such as a user's first name, last name, and logon name). The global catalog contains the information that is necessary to determine the location of any object in the directory.

The global catalog contains a copy of every object from every domain in the forest, but only one set of attributes from each object.

A global catalog server is a domain controller that stores a copy of the global catalog and processes queries to the global catalog. Global catalog servers improve the performance of forest-wide searches in Active Directory. For example, if you search for all the printers in a forest, a global catalog server processes the query against the global catalog and then returns the results. Without a global catalog server, this query requires a search of every domain in the forest.

The first domain controller you create in Active Directory is a global catalog server. You can configure additional domain controllers to be global catalog servers in order to balance the logon authentication traffic and query traffic.

In the directory search user interface, the global catalog is abstracted as the entire directory when selecting a search scope. Users can search the forest without having any prior knowledge of the forest structure. Because there is a single, consistent search interface you don't need to educate users about the directory structure, and you can change the structure within a forest without affecting the way users interact with the directory.

The availability of global catalog servers is crucial to the operation of the directory. For example, a global catalog server must be available when processing a user logon request for a native-mode domain, or when a user logs on with a user principal name.

As a general rule, designate at least one domain controller in each site as a global catalog server.

The following are considerations when placing and creating global catalog servers:

- A computer can start up, but cannot join a domain if a global catalog server is not available and universal groups are involved.
- Cached local logon credentials are all that is available to a user if a global catalog server is not available. This is important for visiting users who are not members of a local domain, but are a member of a remote trusted domain. A global catalog server should be placed at the remote site to provide global catalog access.
- Partial authentication without universal groups is not acceptable because there might be resources for which a user is denied access for universal groups.
- When expanding universal groups, the domain controller contacts a global catalog server, not the local computer at which the user logs on.

- Having too many global catalog servers generates unnecessary replication traffic.
- Domain controllers only perform authentication and the expansion of domain local groups. Global catalog servers expand universal group memberships.
- Global catalog searches offer no benefit in a single-domain environment, but no penalty either. Querying the global catalog produces the identical result of querying a domain controller, because there is only one domain and therefore all domain controllers (even the global catalog servers) hold a full replica.
- The general rule is to have one domain controller per site to serve as a global catalog server to enumerate universal groups.
- The global catalog server provides the address book in Exchange 2000. Therefore, it might be necessary to ensure that a server is local for this service to function.

Operations Master

An operations master is a domain controller that is assigned one or more special roles in an Active Directory domain. The domain controllers that are assigned these special roles perform single-master operations, or operations that are not permitted to occur at different places in the network simultaneously.

Schema Master

The schema master controls all updates and modifications to the schema. To update the schema of a forest, you must have access to the schema master. You can have only one schema master in the entire forest at any time.

Domain Naming Master

The domain naming master controls the addition or removal of domains in the forest. You can only have one domain naming master in the entire forest at any time.

Relative Identifier Master

There is one domain controller that acts as the relative identifier ID master in each domain in the forest. The relative ID master allocates sequences of relative IDs to each of the various domain controllers in its domain.

Whenever a domain controller creates a user, group, or computer object, it assigns the object a unique security identifier (SID). The SID consists of a domain SID that is the same for all SIDs that are created in the domain, and a relative ID that is unique for each SID that is created in the domain.

Primary Domain Controller Emulator

Each domain in the forest must have one domain controller that acts as the primary domain controller (also known as PDC) emulator. If the domain contains computers that are not running Windows 2000 client software, or if it contains domain controllers running Windows NT, the primary domain controller emulator processes password changes and replicated updates to the backup domain controllers running Windows NT.

In a Windows 2000 domain in native mode, the primary domain controller emulator receives preferential replication of password changes that are performed by other domain controllers in the domain. If a password was recently changed, that change takes time to replicate to every domain controller in the domain. If logon authentication fails at another domain controller because of a bad password, that domain controller will forward the authentication request to the primary domain controller emulator before it rejects the logon request.

Infrastructure Master

There must be one infrastructure master in each domain. The infrastructure master is responsible for updating the group-to-user references whenever group memberships change.

If you make modifications to user accounts and group memberships in different domains, there is a delay between the time that you rename a user account and the time that a group that contains that user can display the new name of the user account. The infrastructure master of the group's domain is responsible for this update. It distributes the update through multi-master replication.

Windows 2000 Sites

A Windows 2000 site is a combination of one or more Internet Protocol (IP) subnets that are connected by a high-speed link. Defining sites allows you to configure Active Directory access and replication topology so that Windows 2000 Server uses the most efficient links and schedules for replication and logon traffic.

You create sites for two primary reasons:

- To optimize Active Directory replication traffic.
- To enable users to connect to a domain controller by using a reliable, high-speed connection.

Sites map the physical structure of your network, whereas domains map the logical structure of your organization. The logical and physical structures of Active Directory are independent of each other, which has the following consequences:

- There is no necessary correlation between your network's physical structure and its domain structure.
- Active Directory allows multiple domains in a single site, in addition to multiple sites in a single domain.
- There is no necessary correlation between site and domain namespaces.

Active Directory Schema

The Active Directory schema consists of different objects, or components, that control the classes and attributes maintained by Active Directory. You can add or modify components within the schema, but you cannot delete unused components. Unused schema components can only be deactivated.

> **Note** You cannot deactivate schema objects that are part of the default schema that is included with Active Directory. You can only deactivate schema objects that have been added to the default schema.

Schema Objects

When you create a new user in Active Directory, you create an object of the User class. Do not confuse this with schema modification, which involves creating or modifying the class definitions themselves. Schema modification involves changing the schema components. These components are classes, attributes, and attribute syntax rules.

Classes

Classes are definitions for groups of objects that share a set of characteristics or attributes. For example, User is a class in Active Directory. Every user has certain characteristics in common with other users, such as a first and last name. Although the value of each characteristic is different, they all possess a first and last name.

Each class in Active Directory has a class-schema object corresponding to it in the schema. The class-schema object specifies the attributes and hierarchy rules of the class, and causes the objects created in the class to have the following constraints:

- Objects must contain a list of mandatory attributes that must be present on any object that is an instance of this class.
- Objects might contain a list of attributes that you can place on an object that is an instance of this class.
- Object hierarchy rules are a list of attributes that determines the possible parents in the directory tree of an object that is an instance of the class.

> **Note** An object is only allowed to have an attribute that belongs either to the mustContain or to the mayContain list of classes.

Attributes

Attributes define objects within classes. A sample attribute for an object of the user class might be the user's last name. Each user object will have this attribute, but each will hold a different value that is specific to the user. Every attribute has a corresponding attribute-schema object. The attribute-schema object specifies various properties of an attribute, such as its required syntax and whether or not it can have multiple values.

Syntax Rules

Syntax rules state that attributes can hold specific types of information, such as integer or date-formatted values. For example, only numeric values are acceptable for the *telephoneNumber* attribute. Syntax does not appear as an object in the database; it is hard–coded internally, which means you cannot modify syntax.

Modifying the Active Directory Schema

You can modify the schema directly by using the Microsoft Management Console (MMC) or by scripting, or indirectly by installing software that changes the schema as part of its installation. Software applications that add classes or attributes during the application installation process are referred to as directory-enabled applications. You can change the effect of the schema by deactivating components. It is important to realize the implications of modifying the schema before you make changes.

Using the Active Directory Schema Snap-in

Members of the Schema Admins group can use the Active Directory Schema snap-in to manage the schema by creating, modifying, and deactivating classes and attributes.

You can specify what attributes are indexed and what attributes replicate to the global catalog by using the Active Directory Schema snap-in. For more information about installing the Active Directory Schema snap-in, see Windows 2000 Server Help.

Scripting

You can write a script with Active Directory Service Interfaces (ADSI) that directs the schema to create, modify, or deactivate classes and attributes. You use this method when you want to create schema modifications. Scripting requires that you review the script before running it, to reduce the chance of topological errors.

The following is an ADSI script that adds a user to the schema:

```
Dim oDomain
Dim oUser
Set oDomain=GetObject("LDAP://OU=Test Unit,DC=server,DC=com")
Set oUser = oDomain.Create("user","cn=Test User")
oUser.Put "samAccountName","TUsers"
oUser.Put "givenName","Test"
oUser.Put "sn","User"
oUser.Put "userPrincipalName","tuser@server.com"
oUser.SetInfo
MsgBox "User created " & oUser.Name
Set oDomain = Nothing
MsgBox "Finished"
WScript.Quit
```

Deactivating Schema Components

You can never remove default classes and attributes from the schema; you can only deactivate them. You deactivate or reactivate a class or attribute by using the **Properties** dialog box for that object in the Active Directory Schema snap-in. Before you deactivate a class or attribute, you should consider the following:

- You cannot deactivate default schema objects, only objects that are added after the schema installation.
- You cannot deactivate attributes in the mustContain or mayContain properties of existing active classes.
- When you deactivate a class or attribute, it is no longer replicated throughout the network or to the global catalog server.
- Deactivating a class does not deactivate existing objects of that class type. However, you cannot create new objects from a deactivated class.
- You cannot use deactivated attributes in new or existing classes.
- Objects from deactivated classes and class attributes continue to appear in searches. If you do not need objects of a certain class, you can deactivate the class, search for objects that are instances of the class, and then delete those objects from Active Directory. Similarly, after deactivating an attribute, you can search for objects that have values for that attribute and then delete the attribute's value.
- You cannot create new classes or attributes that have the same common name, object identifier, or Lightweight Directory Access Protocol (LDAP) display name as a deactivated class or attribute.

Implications of Modifying the Schema

Schema modifications impact your entire network. Modifications can affect existing objects, create replication latency, and increase network traffic.

Possible Effects on Existing Objects

A schema update can make an existing object invalid. For example, suppose Widget is an instance of class Products. Products has a mayContain attribute called Color. If a schema update deletes Color from the mayContain list for Products, Widgets becomes invalid, because Widgets has an attribute that is no longer in its class definition.

Active Directory allows deactivated classes and attributes to remain in the directory to ensure that they do not cause any other schema conflicts. However, Active Directory does not remove objects. Therefore you can search on the invalid attribute and delete it from all existing objects. Active Directory does not allow you to add the attribute to any new objects.

Replication Latency

When you perform schema modifications on one domain controller, they replicate across all domain controllers of a newly created class (an instance), and you can replicate them to all domain controllers in the forest. To ensure integrity and expedite convergence, schema replication is a separate process from normal directory replication.

Schema replication can introduce temporary inconsistencies in the schema. It is possible to replicate an object of a newly-created class (an instance) to a domain controller before the new schema class has replicated.

For example, you create a new class called companyVehicles at a domain controller. Then you create an instance of a class, called Limo, at the same domain controller. However, when the changes replicate to another domain controller, Limo replicates before companyVehicles. The replication of the instance fails because the second domain controller is not aware of the class companyVehicles. If a replication failure occurs, Active Directory replicates the schema from the schema operations master, and the schema cache immediately updates on the target domain controller. Active Directory then replicates any object that fails to reach the target domain controller again. As a result, the companyVehicles class is added to the target domain controller schema cache. The next replication cycle successfully adds Limo.

Network Performance During Replication

Some schema modifications can have considerable impact on the network performance. If you deploy Active Directory and then decide to tag another attribute for global catalog replication, you can make the change easily in the schema master. However, this change causes all global catalogs to set their update sequence numbers (USNs) to 0. In Active Directory replication, each Active Directory domain controller maintains a 64-bit counter called an update sequence number. At the start of each update transaction that originates from or replicates to a domain controller, the domain controller increments its current USN and associates this new value with the update request. As a result, all objects in Active Directory (not just the changed property) must be replicated to each global catalog again, which puts a significant load on your network.

Optimizing Exchange 2000 Server

There are many ways that you can optimize your Exchange 2000 organization. The following sections detail how you can optimize your Exchange organization, from using full-text indexing to increase search capabilities for your users to making bulk changes to Active Directory to save time. The optimization methods discussed are as follows:

- Scaling front-end and back-end servers
- Creating multiple virtual servers
- Full-text indexing
- Configuring DNS for a unified namespace
- Accessing Active Directory data
- Making bulk changes to Active Directory
- Storing data in Active Directory
- Preparing for administration

Scaling Front-End and Back-End Servers

You can scale your system to accommodate more users by implementing and configuring front-end or back-end servers and virtual servers.

A front-end server is a server running Exchange 2000 that does not host data, but instead forwards client requests to a back-end server for processing. The front-end server uses LDAP to query Active Directory to determine on which back-end server a user's mailbox resides.

A back-end server is a server running Exchange 2000 that maintains at least one database. This division of features between two servers provides the following benefits in a Microsoft Outlook Web Access environment:

- **Single namespace** You can define a single namespace for users to access their mailboxes. Users can use the same URL, even if servers are added and removed, or if mailboxes are moved from server to server. In addition, creating a single namespace ensures that Outlook Web Access remains scalable as your organization grows.
- **Offload Secure Sockets Layer (SSL)** Encrypting and decrypting message traffic uses processor time. By offloading SLL operations to a front-end server, the back-end server has more processor time to devote to the Microsoft Web Storage System.
- **Firewalls** You can place the back-end server behind a firewall configured to allow only traffic from the front-end server.

For more information about front-end and back-end servers, see Exchange 2000 Help

Creating Multiple Virtual Servers

If you support users with different configuration needs, such as security requirements or message formats, you can create multiple protocol virtual servers. Exchange 2000 allows you to create multiple instances of the protocol server on one computer.

During installation, a default protocol server is created for most protocols. Each of these protocol servers is specifically configured for the protocol used. You can configure items such as authentication methods, message format, and data transfer limits.

You must uniquely identify each virtual server among the other virtual servers for that protocol. To do this, you must specify a unique IP port and address combination for each.

For Hypertext Transfer Protocol (HTTP) virtual servers, you can use the host header parameter to uniquely identify a virtual server.

The following are reasons to create multiple virtual servers:

- **Different Authentication Mechanisms** External users sending messages over the Internet can encrypt all messages with Transport Layer Security (TLS) for additional security. Users on an intranet might not use TLS encryption and do not need to incur the additional cost associated with encrypting the message. Creating two virtual servers, one for encrypted Internet messages, and the other for intranet messages separates these sets of users.
- **Optimize Trusted Applications** Applications that use Collaboration Data Objects (CDO) to send SMTP messages can use an STMP virtual server that is not restricted by reverse Domain Name System (DNS) lookup or recipient limits.
- **Different Server Purpose** The ability to create multiple virtual servers gives you more flexibility when designating a virtual server's purpose. You can connect one virtual server to the Internet, allowing users to send and receive messages over the Internet. And, you can configure another virtual server to only deliver messages within your organization.

For more information about creating and configuring multiple virtual servers, see Exchange 2000 Help.

Full-Text Indexing

Full-text indexing allows users to use Outlook Advanced Find or custom clients to quickly locate messages and documents stored in the Web Storage System that have a certain word either in a message body, message subject, or within an attached document. Users can search other message properties and attributes by using classic search techniques. For more information about enabling full-text indexing, see Exchange 2000 Help.

Full-text indexing searches for words found within messages in addition to words within attachments. It can search only certain types of attachments, which include:

- Embedded Multipurpose Internet Mail Extensions (MIME) messages (.eml)
- HTML (.html, .htm, .asp)
- Microsoft Excel (.xls)
- Microsoft PowerPoint (.ppt)
- Microsoft Word (.doc)
- Text files (.txt)

Searches

The Microsoft Search service builds and uses the full-text index for a store. Searches using full-text indexing are word–based and return the messages that contain the words in the message or the attachment. If a user performs a search for a word and this word is found in the body of a message and in the attachment of another message, the return set will only identify the messages, not the specific attachment within the message. Full-text search supports MAPI and Internet Message Access Protocol version 4 (IMAP4) clients when you enable it. Exchange 2000 Server provides search capabilities on other properties not available in the index.

The Exchange 2000 query processor processes each issued search request. The query processor determines whether Exchange 2000 or Microsoft Search should use the full-text index to process a search. This determination is made based on whether the searched property is indexed. Figure 31.3 illustrates this process.

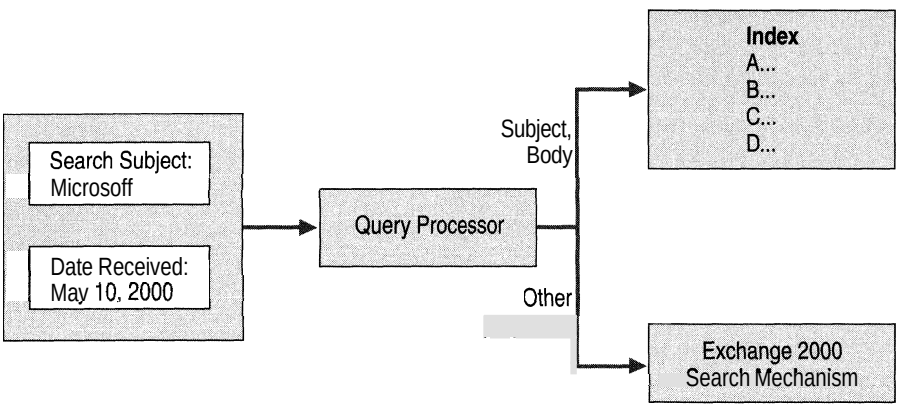


Figure 31.3 The Exchange 2000 query processor

Full-text index searches are only performed when the user specifies a word or words. Because a full-text index does not include all properties of each document in the Web Storage System, there are many circumstances in which a combination of queries must be performed. Subject, Body, To, and From are among the default properties with which a document is indexed. Any other property on which a user searches does not use the full-text index because the index does not include that data.

For example, in Figure 31.3, suppose a user performs a search against an indexed store for documents that contain the word "Microsoft" in the subject received on May 10,2000. The query processor uses Microsoft Search to search the full-text index to locate all documents stored in the database that contain the word Microsoft in the subject. After the full-text search completes, Exchange 2000 searches the result set for all documents received on May 10,2000. The final result set meets the query parameters as set by the user, but requires Exchange to perform two separate searches.

Word-Based

Full-text searches return different sets of data than standard character-based searches. Full-text indexing uses stemming. For example, a search for "comp" will not match with "computer," but a search for "computer" will match with "computers." Stemming will match "computer" with "computers" because they are related words.

Support for Earlier Versions of Client Software

Full-text search is enabled on the server running Exchange 2000 rather than on the client to allow clients running earlier versions of client software to use this feature. As a result, if full-text indexing is enabled on the server, any MAPI client can use it when performing a search of server documents. However, if full-text indexing is not enabled, all MAPI clients perform searches using character-based searches of the message body. These searches are performed at the time of query, not when the index updates.

> **Note** MAPI, IMAP4, and custom clients using WebDAV or Extended OLE DB can take advantage of full-text indexing.

Building the Index

Microsoft Search builds the index and exposes the interface that Exchange 2000 Server uses. Both the Web Storage System and Microsoft Search service must be running for the index to be created, updated, or deleted.

Microsoft Search builds the initial index by processing the entire store, one folder at a time, to identify and log searchable text. By building the index in advance, Exchange can quickly return a search request to the user.

You see heavy CPU use during the indexing process, which can take hours to complete depending on the size of your mailbox store or public folder store. The index occupies roughly 10 to 30 percent or more of the space of the data being indexed, depending on whether cleartext or rich text is being indexed. The index that corresponds to a 5-gigabyte (GB) database occupies about 1 GB of hard disk space.

Changes made in the Web Storage System are added to the index on a schedule. Microsoft Search waits for the scheduled time to update the index to include the new, changed, or deleted items. After the index is created, changes to folders in the Web Storage System will not be picked up until the next automatically–scheduled or manual population (either full or incremental population).

For more information about full-text indexing, see Exchange 2000 Help.

Gather Files

Full-text indexing supports specific file types. If a document is named with an extension indicating a supported file type, but it is not actually that type, Microsoft Search does not index that document. Indexing continues with the next message or document.

A Microsoft Search error is logged in the Windows 2000 Server Application Log at the end of the indexing process. This error records how many documents are not indexed successfully. You can view which documents have problems in the gather file.

Gather files are created during every index process and you can find them in the \Exchsrvr\ExchangeServer\GatherLogs directory. All files with a .gthr file name are text files that you can view to identify every document and message that is not successfully indexed. If there are no files that cannot be indexed, you only see four lines in this file and the file size is less than 300 bytes. Each line in the file after the fourth line identifies the URL of the message or document that fails to index. This line includes the subject or file name along with the error code. The last number that you see in the line is the error code. To decode this error number, use the **gthrlog.vbs** utility found in the \Program Files\Common Files\System\MsSearch\Bin directory. The syntax for this utility is as follows (where *filename* is the name of the .gthr file):

```
cscript gthrlog.vbs <filename>
```

You use the gthrlog.vbs utility from the command prompt and results from the utility appear at the command prompt.

Scheduled Updates

Scheduled updates allow you to update the index during off-peak hours, so that users receive optimum server performance. The disadvantage is that the index can become out–dated between scheduled updates.

Note Scheduled updates should occur at least once a day.

Supporting Clients

If you want to widely implement full-text searching in your organization, consider which messaging client to deploy. MAPI, IMAP4, and custom clients using WebDAV or Extended OLE DB can take advantage of full-text indexing.

Best Practices for Full-Text Indexing

To optimize full-text indexing, you need to prepare your Exchange environment by correctly configuring your server and ensuring your Exchange organization is stable.

Server Configuration

Use a mirrored redundant array of independent disks (RAID) configuration. Microsoft recommends using a RAID 0+1 configuration. This configuration allows for the best performance while ensuring redundancy. A RAID 5 configuration is not recommended.

Disk Space Requirements

The Microsoft Search service requires that the disk that contains the catalog have at least 15 percent free disk space at all times. The size of your catalog will range from 10 percent to 30 percent of the size of your Web Storage System, depending on the types of files stored. It is also important to consider your database growth rate if you plan to retain a large amount of data.

Memory Requirements

Add an additional 256 megabytes (MB) of RAM more than the recommended configuration for an Exchange 2000 server. Microsoft does not recommend running full-text indexing with less than 512 MB of RAM.

File Placement

For optimal performance, place large files and frequently accessed files on separate disks. Microsoft recommends the following:

- Put Exchange .edb and .stm files on a large RAID array.
- Put Exchange log files on the RAID array.
- Put the full-text index on the RAID array. You can specify the index location in System Manager when you create the index.
- Move the full-text index temporary directory to the RAID array (You can use Settemppath.vbs). These files are installed on the system drive by default, which typically does not have the input/output (I/O) throughput of the RAID array.
- A potential improvement is to move the paging file to the RAID array, because it allows faster disk VO by paging before the system exhibits performance problems.

Preparing Your Exchange 2000 organization

Before you install full-text indexing, verify that your Exchange 2000 server or topology is correctly set up and running. Changes made to your Exchange organization after full-text indexing installation could require a full repopulation of the index.

Performance Objects to Monitor

The following objects have counters you should monitor to evaluate full-text indexing:

- Microsoft Gatherer
- Microsoft Gatherer Projects
- Microsoft Search
- Microsoft Search Catalogs
- Microsoft Search Indexer Catalogs

For more information about performance objects and monitoring, see "Monitoring and Maintaining" in this book.

Configuring DNS for a Unified Namespace

Recipient policies are tightly integrated with DNS. To configure a unified namespace, first configure DNS to identify the appropriate computer running Exchange 2000 Server as the mail exchanger (MX) for each e–mail domain you plan to define in your organization. For example, if you have three Exchange servers and each one handles incoming e-mail for multiple departments, and each user has multiple valid SMTP addresses, you must define each computer running Exchange in DNS with an MX record. This will identify the three domains to be handled by a particular Exchange server. Your DNS records might appears as follows for the following zone:

```
        London          A       172.16.1.2
        Tokyo           A       172.16.2.2
        Seattle         A       172.16.3.2
        Europe          MX      10      London
        US              MX      10      Seattle
        JP              MX      10      Tokyo
        Exchange        MX      10      Seattle
        Headquarters    MX      10      London
@                       MX      10      London
@                       MX      20      Seattle
@                       MX      30      Tokyo
```

In this example, all e-mail sent to a user in this zone is delivered to the London server, with backup delivery in Seattle and Tokyo respectively (based on the priority). E-mail sent to user@europe.microsoft.com and user@headquarters.microsoft.com is delivered to the London server. E-mail sent to user@us.microsoft.com and user@exchange.microsoft.com is delivered to the Seattle server, and e-mail sent to user@jp.microsoft.com is delivered to the Tokyo server.

Accessing Active Directory Data

Exchange 2000 accesses Active Directory in many different ways by using many different clients. Users who use Microsoft Outlook 2000 can be assured the most efficient access to the directory. Aside from user access, the Exchange server itself accesses the directory to store critical information. Exchange 2000 speeds up directory access by using a Directory Access Cache.

Outlook 2000

The first time an Outlook 2000 client connects to a computer running Exchange 2000 Server, it looks for the directory service on the home Exchange server. Because the version of Exchange used at the back-end is determined when Emsmdb32.dll loads, Outlook goes through the DSProxy process for the very first session. After the client contacts the DSProxy service (it tries all available transport protocols), a referral is passed back to the client informing it that all future directory requests should go to the global catalog server. Outlook sets the referral in the MAPI profile in the registry:

The referral mechanism reduces load on the Exchange 2000 server and reduces the latency for address book searches. If an explicit server name is entered into the profile, Outlook must restart if that Active Directory server fails. If this occurs, the Exchange 2000 server sends Outlook a new referral.

In some scenarios you might want to force Outlook clients, even the latest versions, to always go through the DSProxy process without being referred. You can do this by configuring the computer running Exchange 2000 to not give out referrals.

To configure an Exchange 2000 server to provide no referrals

1. On the Run line, type **regedt32.exe** or **regedit.exe**, and then click OK.

   In the registry editor, navigate to HKEY_LOCAL_MACHINE\System \CurrentControlSet\Services\MSExchangeSA\Parameters

2. Select the **No** RFR **Service** entry.

3. To prevent an Exchange 2000 server from providing referrals, assign data type DWORD and a value of **0x1**.

   In Regedit.exe, right click the entry, and then click **Modify**

   -or-

   In Regedt32.exe, click the entry, click **Edit,** and then click the appropriate menu item.

4. Close the registry editor.

Compatibility with Earlier Versions of Outlook

If your organization has clients that run Exchange client, Outlook 97, or Outlook 98, you must make accommodations for these clients on your Exchange 2000 server. You must make MAPI Directory Service requests to a server running Exchange.

To make Exchange 2000 compatible with the existing MAPI client base, Exchange 2000-based computers pass any MAPI directory service requests through to a local global catalog server on the network. The directory service proxy (DSProxy) process on the Exchange 2000 server is responsible for this task. Because Active Directory supports a number of protocols, including LDAP and MAPI Directory Service, an Outlook directory request is completely valid, even if it runs directly against Active Directory.

After the global catalog server returns the result to the Exchange 2000 server, the server passes the result to the MAPI client. This process is hidden from the user. The actual communication for one recipient name search is as follows:

1. The MAPI client sends one network packet to the computer running Exchange 2000. The packet contains the name for lookup in plain text.
2. The computer running Exchange 2000 passes the request to a local global catalog server.
3. The local global catalog server returns the request to the Exchange 2000 server.
4. The computer running Exchange 2000 returns the result to the MAPI client.
5. The MAPI client returns an acknowledgement to the Exchange 2000 server
6. The Exchange 2000 server passes the acknowledgement to the local global catalog server.

The directory lookup process produces six frames on the network. The decrease in performance on the global catalog server is between 1 percent and 2 percent. If there are multiple name requests, the name fragments are sent in one request packet.

If the user chooses to browse the global address list, the same process occurs.

Server Directory Access

Depending on the type of query, an Exchange 2000 server might communicate with different Active Directory servers. An Exchange 2000 server usually establishes a number of LDAP connections to domain controllers and global catalog servers that are nearby.

When an Exchange 2000 server needs to look up address information for routing e-mail, for example, to resolve a user name in the directory, the server might choose to forward the request to a local domain controller if the resolution can be processed within the domain. Exchange 2000 uses DNS to find the collection of domain controllers and uses them sequentially. If a local domain controller cannot service the request, it is sent to a global catalog server where it can be resolved. If there is more than one global catalog server in a site, Exchange 2000 uses these on a round-robin

basis. Exchange attempts to connect to up to ten global catalog servers in the same Windows 2000 site as the Exchange server. If all attempts fail, Exchange tries to contact a global catalog server outside of its Windows 2000 site.

Making Bulk Changes to Active Directory

Exchange 2000 Server uses LDAP to access Active Directory. You can use this protocol to make bulk changes to the directory, such as when two organizations merge. Making bulk changes to Active Directory not only optimizes your changes to the directory, but it reduces the time needed to make large changes to the directory.

To make changes to Active Directory using the LDAP Data Interchange Format (LDIF) Directory Exchange Tool, you must write a data file. LDIF is the file format required to import data into Active Directory, and it is also the format in which Active Directory exports data.

> **Note** There are other methods of making changes to the Active Directory. For more information, see the *Windows* 2000 *Sewer Resource Kit Distributed Systems Guide*.

In the following example of an LDIF import file format, you can also see how to add a user object to the myDomain.microsoft.com domain:

```
dn: CN=sampleUser,CN=Users,DC=myDomain,DC=microsoft,DC=com
changetype: add
cn: sampleUser
description: Example of an Imported User using LDIFDE
objectClass: user
SAMAccountName: sampleUser
```

Storing Data in Active Directory

Exchange 2000 does not keep its own directory, but stores data in Active Directory. It is important to understand the data partitions in Active Directory and the global address list, and how to select attributes to replicate to the global catalog, so that you can better optimize your Exchange 2000 organization.

Data Partitions in Active Directory

The information stored in Active Directory on every domain controller in the forest is partitioned into three categories: domain, configuration, and schema. These directory partitions are the units of replication in Active Directory. If the domain controller is also a global catalog server, it holds a partial set of objects stored in the global catalog.

> **Note** You can view the domain, configuration, and schema partitions by using ADSI edit or other LDAP tools, which are included with the Windows 2000 Server Support Tools.

Domain Partition

The domain partition contains all of the objects in the directory for a domain. Domain data in each domain is replicated to every domain controller in the domain, but not beyond its domain. Domain objects include recipient objects such as users, contacts, and groups, and public folder addresses.

Because of the design of the directory structure, the object classes and terms in Exchange 2000 are different from earlier versions of Exchange Server. Table 31.1 compares the object classes and terms between Exchange 2000 and earlier versions of Exchange.

Table 31.1 Object classes and term comparison between Exchange 2000 and earlier versions of Exchange

Exchange 5.*x* Directory Object	Equivalent Object in Active Directory	Comments
Mailbox	Mailbox-enabled user	Mailbox-enabled users are security principals in Active Directory. These users can send and receive messages and have an SMTP address.
Custom Recipient	Mail-enabled contact or mail-enabled user	Mail-enabled contacts are not security principals in Active Directory. All mail-enabled contacts will have an SMTP address. Users on previous messaging systems, such as Lotus cc:Mail and Lotus Notes, are also represented as contacts in Active Directory.
Distribution Lists	Mail-enabled group	Different group classes exist in Active Directory. A group can either be a security or distribution group. Distribution groups are the equivalent of distribution lists in Exchange 5.5.
Public Folder	Public folder	You can only create these object types through the Exchange System Manager.

Note A user object in Active Directory can be mail-enabled only, and not have an Exchange 2000 mailbox. A mail-enabled contact can have an e-mail address that is external to the Exchange organization, but mail-enabled contacts cannot be granted permissions to resources.

Configuration Partition

The configuration of the Exchange 2000 organization is stored in the configuration partition of Active Directory. Because Active Directory replicates the configuration partition between all domains in the forest, the configuration of the Exchange 2000 organization is also replicated to all domain controllers throughout the forest. The configuration partition defines the topology, connectors, protocols, and services settings of the Exchange 2000 organization.

The Exchange 2000 configuration is stored under the following path in the configuration partition:

```
CN=Microsoft Exchange,CN=Services,CN=Configuration,DC=myDomain,
DC=microsoft,DC=com
```

Schema Partition

The schema partition contains all objects types (and their attributes) that can be created in Active Directory. This data is common to all domains in the domain tree or forest, and is replicated by Active Directory to all domain controllers in the forest.

During the installation of the first computer running Exchange 2000 Server in the forest, the Active Directory Schema is extended with new attributes for Exchange 2000 that start with ms-Exch. The schema is extended using LDIF files.

> **Note** You can examine which attributes are added to Active Directory by viewing the LDIF files on the Exchange 2000 CD.

Global Address List

When Outlook 2000 users look up e–mail addresses for other users in their Exchange organization, the information comes from the global address list. The global address list represents an aggregation of all messaging recipients in the enterprise. Because computers running Exchange 2000 do not host their own directory service, you retrieve all data from the global catalog servers in Active Directory. A global catalog server can support the MAPI protocol in addition to LDAP, so Outlook clients can communicate with Active Directory by using the same protocol used by the directory service in Exchange Server 5.5. By using the Address List Manager, you can create and manage address lists. For more information, see Exchange 2000 Help.

Selecting Attributes to Replicate to the Global Catalog

The global catalog holds a partial replica of domain data directory partitions for all domains in the forest. By default, the partial set of attributes stored in the global catalog includes those attributes most frequently used in search operations, because one of the primary functions of the global catalog is to support clients querying the directory.

> **Note** You can select which attributes to replicate in the global catalog by using the Active Directory Schema snap-in. The Active Directory Schema snap-in is part of Windows 2000 Server Administration Tools, which is included on the Windows 2000 Server compact disc set.

Selecting which attributes to replicate to the global catalog requires careful planning. You need to preserve features that users of Outlook already have if an earlier version of Exchange is deployed, but you have to take into consideration the effect for replication traffic if you select too many additional attributes. Also, marking existing attributes to replicate to the global catalog causes replication of all global catalogs. This can cause replication and network traffic issues for organizations with large global catalogs.

Because the global catalog holds a partial replica of its home domain and a partial replica of every other domain in the forest, users see all attributes for other users in the same domain. However, they see only the attributes tagged for replication in the global catalog from other domains.

If your network is very slow, you might want to survey your Outlook users to find out which directory attributes they rely upon. It is very important to establish whether any custom CDO or ADSI applications rely on the presence of certain directory data. For example, a workflow application might require access to a custom attribute that holds a manager's approval limit.

Each additional attribute tagged for replication incurs an additional replication data per object. You might need to reduce the number of attributes that are tagged for replication due to bandwidth constraints.

Replication in Exchange 5.5 or Earlier

The replication traffic caused by an existing Exchange 5.5 or earlier network is far greater than the traffic produced by Active Directory. This is because each computer running Exchange Server 5.5 in an organization hosts a full copy of the Exchange directory, whereas Active Directory only replicates to domain controllers and global catalog servers.

Any change to an Exchange Server 5.5 object causes the entire object to replicate again to the rest of the Exchange organization (consuming roughly 5 kilobytes (KB) within a site, and 1 KB between sites), whereas Active Directory uses per-property replication, so the amount of replication data is much smaller, and compression is better.

Preparing for Administration

Microsoft Exchange 2000 Server is typically administered from the administrator's computer rather than directly from the server. The administrator's computer is most likely a computer running Microsoft Windows 2000 Professional, but you can also use the Terminal Services client, running on Microsoft Windows 95 or Microsoft Windows 98, to access the server.

Configuring an Administrator Computer

You administer Exchange 2000 users solely by using the Active Directory Users and Computers snap-in. You perform all user administration from an organizational unit container such as Users, which is available under the domain name. You use Exchange System Manager to configure Exchange 2000 system settings and servers.

Note You should create custom organizational units to contain your user accounts.

You can install System Manager on any computer running Windows 2000 Server.

To configure an administrator computer

1. In Control Panel, double-click **Add/Remove Programs,** and then click **Add/Remove Windows Components.**
2. Select **Internet Information Services,** and then click **Details.**
3. Click **Simple Mail Transfer Protocol (SMTP)** to add it.

   Note The SMTP service is required for Exchange 2000. You cannot add the SMTP service during Windows 2000 Professional installation.
4. In **Add/Remove Programs,** click **Add New Programs,** and then click **CD or Floppy.**
5. On the Windows 2000 Server compact disc, open the **I386** folder.
6. Click **Adminpak.msi** to install it.
7. Run Exchange 2000 Setup and install only the System Management tools

   Note Windows 2000 Server does not support remote administration of Network News Transfer Protocol (NNTP). You cannot view the current SMTP sessions from a remote computer. You can perform all other Exchange 2000 administration from a remote computer. The Active Directory Users and Computers snap-in is installed on computers hosting Active Directory or Exchange 2000 Server.

Administering Exchange 2000 and Exchange Server 5.5

You can install Exchange 2000 administrative tools on the same computer that is running the Exchange Server 5.5 administrative tools. However, you need to uninstall the Exchange Server 5.5 tools first so that Exchange 2000 does not detect a previous installation of Exchange Server 5.5 and prevent you from continuing with the installation (because Exchange 2000 does not support upgrades). You can then install the Exchange Server 5.5 Administrator program with Exchange 2000.

Tuning Exchange 2000 Performance

Exchange 2000 performance tuning involves making performance decisions and setting registry keys. The following sections illustrate ways to tune your Exchange 2000 organization.

- DSAccess Settings
- Relocation of Database, Log, and .stm files

DSAccess Settings

In many organizations, the global catalog server is an important resource that multiple applications need to access. If there are only one or two global catalog servers in a physical location, monitoring its workload is especially important. To reduce the burden on the global catalog, each Exchange 2000 Server has a DSAccess cache.

The DSAccess cache allows Exchange services to cache directory lookups without querying the global catalog. All directory access, apart from address book searches from MAPI clients and certain portions of SMTP inbound or outbound routing, goes through DSAccess and the DSAccess Cache. This feature increases the performance of the network and the servers running Active Directory.

By default, up to 4 MB of directory entries are cached for a period of five minutes. You can monitor the effectiveness of the cache by using the Windows 2000 Server System Monitor.

When tuning a caching mechanism for a particular environment, it is important that you achieve the correct balance between performance and data access. For example, it is useful to increase the amount of time entries can exist within the cache to further reduce network traffic. However, the cache can become old if entries remain there for too long, whereas increasing the total time that entries can remain in the cache also increases the amount of system memory the cache uses. You must carefully consider the environment before making adjustments, and you need to thoroughly test and document any changes.

DSAccess Cache

The number of users, the maximum number of entries, the maximum cache size (memory) and the cache expiration time are all parameters that can affect the optimal size and performance of the DSAccess cache. The following registry keys provide you with low-level control of the cache.

Cache Expiration Time

The following registry key sets the Time to Live setting (CacheTTL) for entries in the cache. In most situations, the Time to Live setting is governed by the user. This registry key is essential to determining the overall size of the DSAccess cache.

To set the Time to Live setting

Caution Do not use a registry editor to edit the registry unless you have no alternative. Registry editors bypass the standard safeguards provided by administrative tools. These safeguards prevent you from entering conflicting settings or settings that are likely to degrade performance or damage your system. Editing the registry directly can have serious, unexpected consequences that can prevent the system from starting and require that you reinstall Exchange 2000 Server or Windows 2000 Server. To configure or customize Windows 2000, use the programs in Control Panel or Microsoft Management Console (MMC) whenever possible.

1. On the **Run** line, type **regedt32.exe** or **regedit.exe** and the click **OK.**
2. In the registry editor, navigate to HKEY_LOCAL_MACHINE\System \CurrentControlSet\Service\MSExchangeDSAccess\Instance0

   Note Create the **Instance0** subkey if it does not already exist.
3. Select or create the **CacheTTL** entry.
4. To set the Time to Live setting, assign data type **REG_DWORD** and value **0x600.**

   In Regedit.exe, right-click the entry and then click **Modify.**

   -or-

   In Rededt32.exe, click the entry, click **Edit,** and then select the appropriate item.
5. Close the registry editor.

The following two registry keys, maximum number of entries (MaxEntries) and maximum cache size (MaxMemory) are two ways that you can control the size of the DSAccess cache. Properly configuring the cache size by using either registry key involves a certain amount of trial and error. A long Time to Live setting for the cache expiration time parameter can result in increased memory consumption because entries in the cache are retained longer.

Note It is recommended that you configure the cache size using the MaxMemory registry key. The advantage of using this key, as opposed to MaxEntries, is that you have a higher degree of control over the memory usage of the server.

Maximum Number of Entries

Each client (such as Post Office Protocol version 3 [POP3] authentication and SMTP inbound) can involve multiple cache entries. The MaxEntries registry key sets the maximum number of entries in the DSAccess cache.

To properly set this registry key, you should first determine what user actions the server supports. Next, you should determine the number of entries that correspond to each of these user actions. If multiple user actions or protocols are supported, it is recommended that the maximum number of entries that corresponds to a particular user action serve as the determining factor for this registry key. Finally, you set the registry key value to the product of the number of actions multiplied by the number of entries per action, plus the number of client actions you expect to cache at any given time.

To set the maximum number of entries in the DSAccess cache

1. In a registry editor, navigate to HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\MSExchangeDSAccess\Instance0

   Note Create the **Instance0** subkey if it does not already exist.

2. Select or create the **MaxEntries** entry.
3. To set the maximum number of entries, assign data type REG–DWORD and a value for the number of entries. The default value, **0,** allows unlimited entries.
4. Close the registry editor.

Maximum Cache Size (Memory)

Each client action (such as POP3 authentication and SMTP inbound) can involve multiple cache entries. Most client actions consume on average 3.6 KB of cache memory. There is also 2.5 MB of overhead with DSAccess. You can use the following formula to configure the cache size of DSAccess:

```
MaxMemory = 3.6 * (Load Rate) * (CacheTTL) + 2500
```

If, for example, you set the Time to Live setting to the default value of 600 seconds, the server is expected to handle 20 client actions per second.

```
MaxMemory = 3.6 * 20 * 600 + 2500 = 45,700 Kb or 45.7 MB.
```

To set the maximum cache size

1. In a registry editor, navigate to HKEY_LOCAL_MACHINE\System\CurrentControlSet \Services\MSExchangeDSAccess\Instance0.

   Note Create the **Instance0** subkey if it does not already exist.

2. Select or create the **MaxMemory** entry.
3. To set the maximum cache size in memory, assign data type REG–DWORD and a value for the number of kilobytes. The default value is **4096.**
4. Close the registry editor.

DSAccess Configuration

The ConfigDCHostName and ConfigDCPortNumber registry keys tell DSAccess which domain controller to use to search for configuration information.

Config DC Host Name

The following registry key indicates which domain controller Exchange 2000 Server should use to search for configuration information.

To specify a domain controller for searching configuration information

1. In a registry editor, navigate to HKEY_LOCAL_MACHINE\System\CurrentControlSet \Services\MSExchangeDSAccess\Instance0.

   Note Create the **Instance0** subkey if it does not already exist.

2. Select or create the **ConfigDCHostName** entry.
3. To specify a domain controller, assign data type REG–SZ and the value to the name of the domain controller (for example, *seattle.microsoft.com*).
4. Close the registry editor.

Config DC Port Number

The following registry key indicates which port number Exchange 2000 Server should use when communicating with a given domain controller to search for configuration information.

To specify a port for searching for configuration information

1. In a registry editor, navigate to HKEY_LOCAL_MACHINE\System \CurrentControlSet\Services\MSExchangeDSAccess\Instance0.

   Note Create the Instance0 subkey if it does not already exist.

2. Select or create the **ConfigDCPortNumber** entry.
3. To set a port number for searching, assign data type REG–DWORD and a value for the port number (for example, 0x389).
4. Close the registry editor.

Relocation of Database, Log, and STM Files

On most hardware platforms, there is a substantial performance advantage to separating the log files that write sequentially to the hard disk from the database and .stm files that perform random I/O. Even if disk I/O is not a significant factor, separating these files helps analyze the I/O burden of the specific user actions by isolating different features. On systems with multiple hard drives, an ideal configuration is as follows:

- One spindle (physical drive) for binaries
- One spindle for the system swap file
- One spindle dedicated to each storage group
- Separate striped sets containing as many spindles as possible for each .edb and .stm file

At a minimum, you should separate random and sequential I/O so that logs and databases are on separate spindles. It is recommended that you have some type of parity or mirroring such as RAID 1 or RAID 5 in a production environment. This secures your data. When you are testing performance, RAID 0 is recommended.

For more information about moving databases, see Exchange 2000 Help.

Optimizing Message Transport

Exchange 2000 Server builds on the Exchange Server 5.5 message transfer agent (MTA) but uses a full-featured SMTP transport for all native communications. The X.400 protocol supports additional features, such as Request for Comments (RFC) 2156 Multipurpose Internet Mail Extensions (MIME) Internet X.400 Enhanced Relay interoperability. However, in many circumstances, you use the X.400 MTA to connect Exchange 2000 servers with other external X.400 systems rather than for native message transfer between Exchange 2000 servers.

Using SMTP as the native communication method between Exchange 2000 servers eliminates some of the deployment issues of earlier transport implementations. For example, organizations with a distributed user base typically design their Exchange 5.5 site models based on available network bandwidth instead of designing them for convenience of administration. This is because all Exchange 5.5 servers within a site use remote procedure calls (RPC) to communicate with one another, and low-bandwidth and high-latency networks are inefficient for the synchronous nature of RPC. Because Exchange 2000 Server doesn't use RPC to transfer messages, you can create a more flexible routing topology even when low-bandwidth and high-latency networks are unavoidable. The following sections illustrate ways to optimize your Exchange 2000 organization's messaging:

- Overview of message routing in Exchange 2000
- Message routing and group expansion
- Connecting routing groups
- Link state table
- Message routing
- Multiple public folder trees
- Managing public folders

Overview of Message Routing in Exchange 2000

The routing process for a message begins when the Exchange 2000 server receives a message either from a user, another Exchange server, or an external messaging system through a connector or gateway. The Exchange 2000 server looks up the recipient in Active Directory and routes the message to the recipient.

For more information about message routing concepts, see Exchange 2000 Help.

Routing Groups

Routing groups are groups of Exchange 2000 servers that are connected over reliable, permanent links. From a topological perspective. routing groups are equivalent to sites in earlier versions of Exchange. Separating servers into routing groups enables you to control e-mail flow, troubleshoot message transfer between groups of servers, and track messages.

Exchange 2000 server uses SMTP as its primary protocol, rather than RPC. With SMTP, it is not as important to have a high bandwidth connection between servers. SMTP is more reliable than RPC over limited bandwidth connections. It is more important for servers in a routing group to have permanent, reliable connectivity than it is for them to have high bandwidth between them.

Connectors

Routing groups are connected using connectors. The following is a list of connectors available with Exchange 2000 Server that can connect routing groups:

- Routing Group connector

  Simple Mail Transfer Protocol (SMTP) connector

- X.400 connector

You cannot use connectors designed for third party e-mail systems to connect routing groups.

Link State

The Exchange 2000 server routing and selection process uses a link state table to determine the shortest path between two routing groups from a given message. The link state table is stored on each Exchange 2000 server and contains the status of each connector in the Exchange 2000 organization. If a server running Exchange 2000 cannot find a route for a particular message after referring to the link state table, it does not attempt to deliver the message.

Message Routing and Group Expansion

All message routing information, including routing groups and bridgehead servers, is held in the configuration naming context of Active Directory. To make a routing decision, the Exchange 2000 Server contacts a local domain controller and retrieves this information.

If a message is sent to a universal group, the SMTP virtual server, which is configured to perform the expansion, uses LDAP to contact a global catalog and populate the message header with the group membership. If the message is for a domain local or global group, the expansion server should be in the same domain as that group and should be configured to use only global catalogs from the local domain where the group resides. By default, every Exchange server will try to use global catalogs from their local domain and site; however, if there are not enough global catalogs in the local domain and site, Exchange will broaden the scope and request global catalogs just from the local site. Exchange 2000 will preferentially use the global catalogs that are closest to it; however, if they become unavailable, a global catalog that is not from the same domain as a domain local or global group will be used and this will result in e-mail not being sent to the membership of the domain local or global group. Under these circumstances, Exchange 2000 may use up to 10 or more global catalogs from the local domain in the same site as the expansion server. Alternatively, there should be no global catalogs from any other domain in that site, or DSAccess should be configured to use only global catalogs from the same domain as the expansion server and the global or domain local group.

When to Use Multiple Routing Groups

If you have a single physical location and all servers are connected through a reliable, permanent link, you might not need multiple routing groups. However, you might need multiple routing groups under one or more of the following conditions:

- Network connectivity is unreliable.
- You want to control the message paths in an Exchange organization; for example, if you need to alter ihe messaging path from a single-hop to multi-hop, such as when servers are located in separate physical locations, but the servers are configured to communicate in a single-hop environment.
- You want to schedule messaging between two locations.
- You want to control public folder referrals. Public folder referrals will preferentially go to a server in the same routing group, and then through the least cost route between routing groups by using connectors that allow referrals.

Routing Group Topology

The link state provides flexibility in designing routing groups, particularly by allowing multiple paths between routing groups.

Multiple Paths Available Between Routing Groups

In earlier versions of Exchange, the site topology takes into account the possible message bouncing that occurs when sites have multiple routes between them. This means that most site topologies are hub-and-spoke, with only a single connector between the hub and each spoke. Exchange 2000 uses the link state table, which makes this topology unnecessary because messages do not bounce back and forth. If a connector fails, the message can be rerouted through another routing group, and will bounce back to the original routing group because Exchange determines if the connector is down and propagates that state information around the organization.

Message Traffic Analysis Can Help Define Boundaries

Message traffic analysis can provide helpful information about which servers most often communicate with each other. You might consider putting servers that communicate regularly into the same routing group. However, you should not design routing groups based solely on traffic analysis without looking at management overhead, because this might not be the most efficient solution.

Additional Considerations

Some additional considerations can impact how you design a routing group:

- Servers might connect over slow (but reliable) links as part of the same routing group.
- Routing groups are dynamic, and they can change at any time.
- The routing group architecture determines public folder access.

Connecting Routing Groups

Servers in different routing groups communicate using connectors. The following sections detail how to connect your routing groups:

- Planning routing group boundaries
- Routing group deployment example
- Routing Group connector
- SMTP connector
- X.400 connector

The Routing Group connector can only connect routing groups. Other connectors provide connections to external systems in addition to connecting routing groups within the Exchange 2000 organization.

When you configure a connector on a server, the connector is included in the routing process for messages destined outside of the routing group. Servers running Exchange 2000 that host routing group connectors are called bridgehead servers. All messages that are delivered through routing groups pass through the bridgehead server that hosts the routing group connector.

Planning Routing Group Boundaries

The following are prerequisites when grouping Exchange 2000 servers into a routing group:

- Permanent, reliable network connections
- Contact to the routing group master

The following are reasons to divide Exchange 2000 servers into multiple routing groups:

- Inconsistent routing group prerequisites
- Unreliable network connections
- The messaging path must be altered from a single-hop to a multi-hop
- Messages must be queued and sent on a schedule
- Low-bandwidth network connections exist for which X.400 connectivity is more appropriate
- Client connections to public folders

Routing Group Deployment Example

Figure 31.4 shows the physical location of users and network links. The geographic locations are connected with network links of varying speeds. Although there appears to be a high-bandwidth link between two locations, you should consider other traffic that might be present on the line when calculating the available bandwidth.

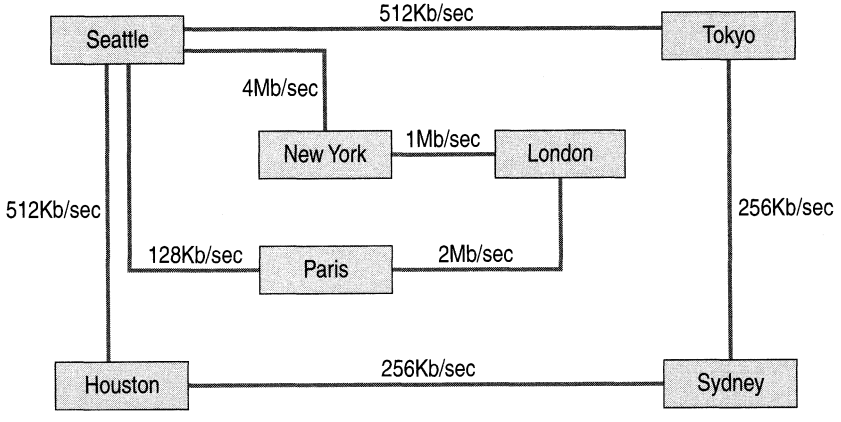


Figure 31.4 Physical location of users and network links deployed

When calculating available bandwidth, you should consider the following variables:

- Whether there are certain groups of users who send large messages to one another on a regular basis.
- What the average size is of messages that travels across the network.
- Which public folders users use.

Figure 31.5 shows the routing group information that you can place over the network structure that results in optimal message transport.

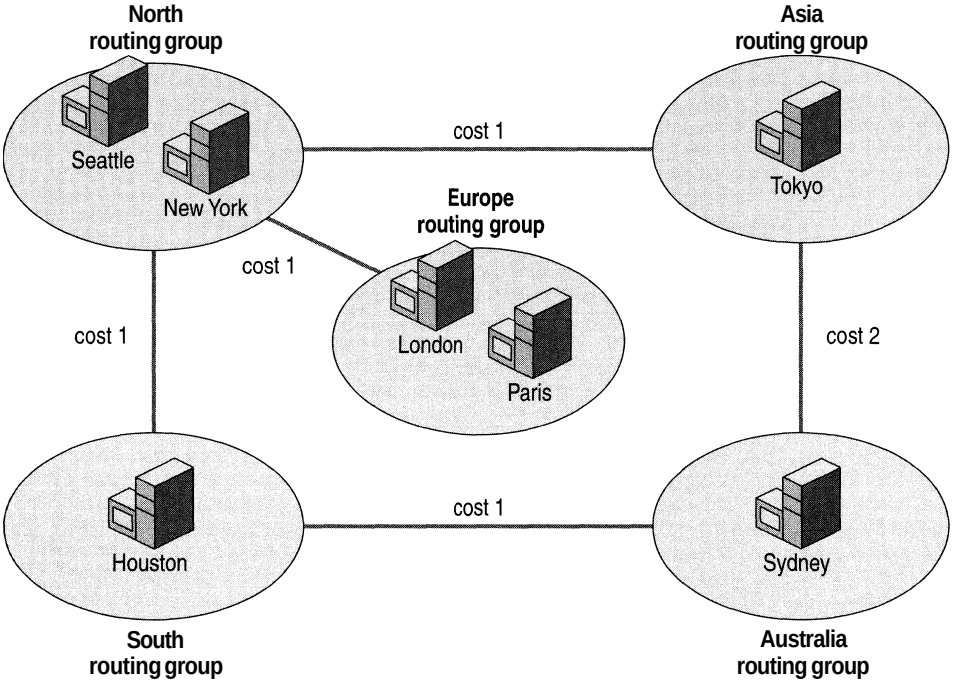


Figure 31.5 Final routing group strategy

As you can see from the final strategy, you can group multiple locations together to form single routing groups, although you must understand the impact this has on client connections to public folders.

This example uses Routing Group connectors because they offer the best features and resilience. Each routing group connector uses the concept of source and target bridgehead servers. To get the maximum efficiency from the network, you can configure bridgehead servers so that all connections occur over a single network link. For example, the target bridgehead servers specified in the connector from the South routing group to the North routing group include connector servers only in Seattle, because no direct network connectivity exists between Houston and New York.

Unlike the Site Connector in Exchange 5.5, multiple target bridgehead servers specified on a connector are not used in a cost-weighted mechanism, nor are they used sequentially. When message transfer takes place, a local bridgehead server reads the list of target servers and chooses the first one on the list. If the first server is down, the bridgehead server uses the second server, and so on. Subsequent messages use the same algorithm.

Routing Group Connector

The Routing Group connector provides the simplest, easiest way to connect two routing groups. The Routing Group connector is similar in function to the Site Connector in earlier version of Exchange. However, the Routing Group connector uses SMTP message transport protocol rather than remote procedure calls (RPC) to deliver messages. Benefits of the Routing Group connector include the following:

- It is easy to configure.
- You can configure it with multiple target bridgehead servers.
- You can use it to connect to earlier versions of Exchange that are configured with the Site Connector, which uses RPCs.

Multiple Bridgehead Servers

You can configure the Routing Group connector with one or more bridgehead servers on either end of the connector. This allows you to control which servers send and receive messages between routing groups. An advantage of having multiple servers identified as bridgehead servers for a Routing Group connector is that if a bridgehead server is not functioning, Exchange can choose another bridgehead server within the routing group to transmit the message.

Messaging Security

Exchange 2000 Server provides SMTP authentication between bridgehead servers. Message encryption between bridgehead servers is disabled by default. If encryption is required, you can implement Internet Protocol security (IPSec), which is a standard encryption method for TCP/IP network security. For more information about IPSec, see "Security" in this book.

Resolving the Target Sewer IP Address

Exchange 2000 server performs mail exchange (MX) record resolution for SMTP but usually resolves the target server by using an A (host) record. This simplifies the configuration of the Routing Group connector because an Exchange or DNS administrator does not have to create or manage MX records.

When a bridgehead server that hosts the Routing Group connector receives a message to deliver across the connector, the bridgehead server tries to resolve the target server's IP address by using the standard SMTP resolution process. That is, the bridgehead server first tries to resolve the target server defined on the Routing Group connector by using DNS MX records. If no MX records exist for the target server, it performs a DNS query for an A record for the target server. This means that an A record must exist in DNS for all servers running Exchange. The Windows 2000 DNS service registers A records for all servers running Windows 2000 Server, including all Exchange 2000 servers.

If for some reason, an A record for the target server is not found, the bridgehead server tries to resolve the IP address by using the network basic input/output system (NetBIOS) name resolution process.

When multiple target servers exist for a Routing Group connector, Exchange 2000 server intercepts the request, looks up the bridgehead servers, and returns them to SMTP before SMTP resolves the MX record against a DNS server.

SMTP Connector

You can use SMTP connectors to connect routing groups. You should use them when the following conditions exist:

- When the remote server connector is the Internet Mail Service from earlier versions of Exchange.
- When a pull relationship is required between servers in which one side queues messages and the other side pulls them by using the TRN or ETRN command.
- When you want to define Transport Layer Security (TLS) or other security parameters.

Bridgehead Servers

Each SMTP connector can specify multiple bridgehead servers within each routing group. All messages from the bridgehead servers to the target routing group are either delivered directly after MX resolution by DNS or are forwarded to a smart host.

Resolving the Destination Sewer IP Address

The SMTP connector uses the standard method of resolving the destination server's IP address through DNS MX records.

The SMTP connector first tries to resolve the destination server by using DNS MX records. If an MX record does not exist, the sending server attempts to resolve the destination server by using the host name resolution process, which includes querying DNS for an A record. If it still does not find the destination server, the bridgehead resolves the IP address by using the NetBIOS name resolution process.

When multiple servers exist for an SMTP connector, Exchange 2000 intercepts the request, looks up the target servers, and returns them to SMTP before SMTP resolves the MX record against a DNS server.

With an SMTP connector, you can optimize the message transport to a greater extent than with the Routing Group connector. SMTP connector options include authenticating remote domains before sending e-mail, designating specific times when e–mail can send, and setting multiple permissions levels for multiple users on the connector.

Using a Smart Host or DNS Resolution of MX Record

A smart host is an intermediary host that uses DNS to resolve the destination host's IP address and then sends the message to the destination host. The smart host server must be able to process e-mail for the remote address space or routing group that the SMTP connector needs to reach. The SMTP connector relays all messages through the smart host, which passes them on to the remote destination by using DNS.

A smart host is helpful for messages traveling between servers over the Internet, such as when the remote domain can only be reached during certain times or infrequently. Instead of repeatedly contacting the domain until a connection is made, the server running Exchange only needs to transit to the smart host. Then the smart host makes the remote connection.

If a smart host is not designated, a DNS lookup is made on every address to which the SMTP connector sends email.

> **Note** If you use an IP address to identify the smart host, enclose the address in brackets ([]) to increase system performance. The SMTP service first checks for a server name, then an IP address. The brackets identify the value as an IP address, so that it bypasses the DNS lookup.

Remote Triggered Delivery

You can configure the SMTP connector to retrieve queued email from a remote SMTP server at specified intervals. That is, you can configure a remote domain to receive and hold email on behalf of the destination domain. Messages sent to the remote domain are held until the SMTP ETRN command is received from an authorized account on your local Exchange 2000 server. You can also use the ETRN command to remove email from a queue. You select **Request ETRN/TURN when sending messages,** then select the times you want the SMTP connector to contact the remote domain and trigger the delivery of queued e-mail.

Link State Table

Exchange 2000 determines the route that messages will take based on a least cost algorithm. However, a server running Exchange 2000 also has a map of the entire message topology of which it is a member. This map, represented in the link state table, is updated regularly and propagated among all the servers so that each server can determine not only the cheapest way to deliver a message, it can also determine whether all of the connectors that make up the route are functioning. The available routes and costs in the organization are available to any Exchange 2000 server.

Link State Information

There are only two states for any given link in an Exchange 2000 organization: up or down. Exchange 2000 does not propagate connection information, such as whether a link is active or in a retry state. This information is only known on the server involved in the message transfer.

Routing Group Masters

You designate a server to be the routing group master for a routing group. The routing group master maintains link state information received from different sources. The routing group master tracks this data and propagates it to the rest of the servers within the routing group. The first server added to a routing group becomes the routing group master. However, you can change this by using the Exchange System Manager.

If the routing group master fails, you must designate a new routing group master. The servers running Exchange 2000 in the routing group continue to use the existing link state table, without updates, until the original routing group master comes back online or a new routing group master is designated.

> **Note** You can view the status of connectors by using System Manager, click **Tools,** and then click **Monitoring and Status.**

Link State Algorithm

Although Exchange 2000 uses routes and costs, it also uses a link propagation protocol, called the link state algorithm. The link state algorithm is responsible for propagating the state of the messaging system in almost real-time to all Exchange 2000 servers. This has the following advantages:

- Each server running Exchange makes the best routing choices at the time it receives a message by using the latest information from the link state table, so that messages are not sent along a path where there is a failed link.
- The link state algorithm eliminates message bounce between servers, because each server running Exchange 2000 knows if other alternate or redundant links are up or down.

Link State Table Maintenance

If a connection fails between two routing groups, Exchange updates the link state table and notifies all servers in the organization.

Link State Propagation

Link state data is propagated between routing groups through SMTP on port *25*, and within the routing group by using TCP port 691. If connection status changes, the link state table is updated as follows:

1. The bridgehead server with the new connection status marks the connector as up or down in the link state table.
2. The bridgehead server updates the routing group master over TCP port 691
3. The routing group master updates its link state table and updates all of the other servers running Exchange 2000 in the routing group over TCP port 691.
4. All other bridgehead servers in the routing group update the routing groups to which they are connected over TCP port 25.
5. Those bridgehead servers update their routing group masters over TCP port 691, and so on until all routing groups are updated.

Exchange 2000 Server uses an SMTP command rather than a message to transfer link state status.

Recovering a Connection

When a connection fails, the bridgehead server with the failed connector continues to retry the connection three times at 60-second intervals, and then retries according to the schedule set on the **Delivery** tab of the SMTP virtual server.

Although messages are rerouted, the bridgehead server continues to try to open port *25* on the destination server. After a connection is re-established, the bridgehead server notifies the local routing group master that the connection is available. The routing group master, in turn, notifies all of the servers in the routing group that the connection is available. Finally, all bridgehead servers in the routing group notify the bridgehead servers to which they are connected in adjacent routing groups.

Message Routing

Message routing is the process of transferring messages between servers, routing groups, and over the Internet. The following are examples of message routing between routing groups and to various types of message routing topologies.

Routing Messages Within the Same Server

When the server running Exchange 2000 determines that a recipient is on the same server as a sender, Exchange 2000 delivers the message to the recipient's mailbox in the following way:

1. A MAPI client sends a message.
2. The message goes to the advanced queuing engine whether the recipient is a local or remote user.
3. The advanced queuing engine places the messages in the pre-categorizer queue.
4. The message categorizer retrieves the message from in the pre-categorizer queue and processes the message. The message categorizer expands groups when appropriate and checks sender and recipient limits.
5. If the recipient is local, the message categorizer places the message in the local delivery queue.
6. Web Storage System associates the message with the recipient's mailbox.

Routing Messages Within the Same Routing Group

When the server determines that the recipient's mailbox is on a different server within the same routing group, the routing process is slightly more complicated than when messages are routed within the same server. If the message recipient's server runs Exchange 2000, the message is routed through SMTP to the recipient's server. If the message recipient's server is running an earlier version of Exchange, the message is routed to the recipient's servers through RPC. Regardless of the protocol used, message transfer within a routing group is point–to–point, meaning that the originating server communicates directly with the recipient's server. Figure 31.6 illustrates point-to-point message routing:

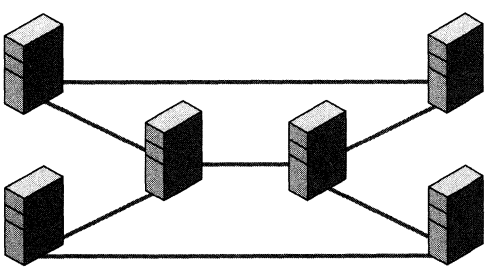


Figure 31.6 Point-to-point routing

The Web Storage System running on the sender's server routes the message though SMTP to the recipient's server in the same routing group. The recipient's server receives the message and delivers the message to the recipient's mailbox.

For more information about the steps involved in message routing, see "Exchange 2000 Architecture" in this book.

Routing Messages to a Server in a Different Routing Group

When the recipient's mailbox is on a sewer in a different routing group, the message must be transferred over a routing group connector. The sender's server identifies a route for the message to take and routes the message to the appropriate bridgehead server. The bridgehead server sends the message to the bridgehead server in the recipient's routing group. The receiving bridgehead server routes the message to the recipient's server.

This method of routing is also referred to as ***store*** and forward. Figure 31.7 illustrates store and forward message routing.

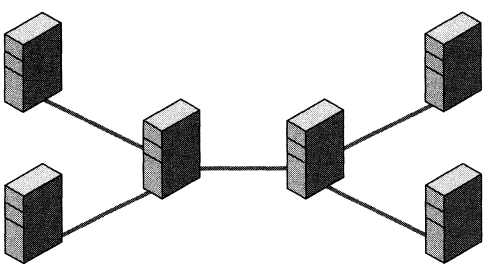


Figure 31.7 Store and forward message routing

Determining a Route Through Multiple Routing Groups

If multiple routing groups exist within an organization, the sending server uses the link state table to determine the best route based on connection cost and status. The message then goes through the appropriate bridgehead servers until it reaches the destination routing group.

Each bridgehead server repeats the routing and selection process by referencing the link state table and routes the message to the next bridgehead server. This process repeats until the message arrives at the recipient's routing group. Depending on the path a message must take to arrive at the recipient's mailbox, the message might travel through multiple servers and routing groups.

If multiple routes have the same cost, the server chooses a random route, to provide load balancing.

In Figure 31.8, a user in the South routing group sends a message to a recipient in the Asia routing group. The least cost route between the South routing group and the Asia routing group is through the North routing group.

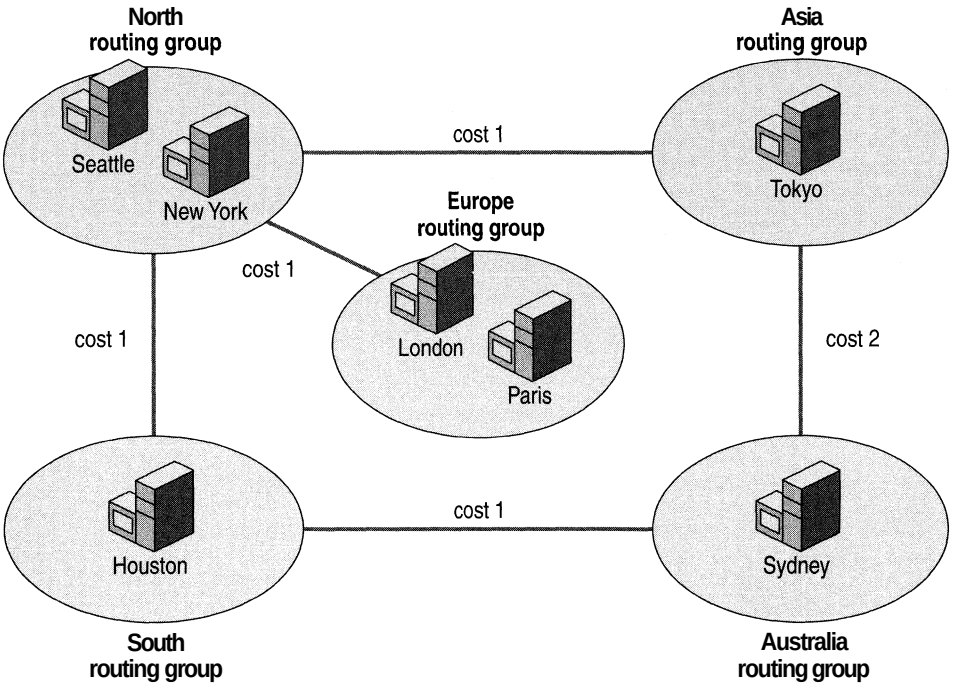


Figure 31.8 Message routing through multiple routing groups

The route the message takes is as follows:

1. The message is sent from the sending user's server to the bridgehead server in the South routing group that connects to the bridgehead server in the North routing group.
2. The bridgehead server in the North routing group receives the message, and sends the message to the Asia routing group through its bridgehead server connected to the Asia routing group.
3. The bridgehead server in the Asia routing group receives the message and sends it to the recipient's server.

Routing Messages Outside an Exchange 2000 Organization

Routing group connectors operate only within an Exchange organization. If a message needs to be transferred to a recipient whose mailbox resides on another messaging system or within another Exchange organization, the message must transfer over a connector that connects to that other system. The message might need to travel through one or more routing groups before finding a bridgehead server that hosts a connector to that messaging system. It is not necessary for every routing group in an Exchange organization to host connectors to other systems.

Rerouting Mail

If a connection fails between routing groups, messages are rerouted. For example, in Figure 31.9, if a user sends a message from the South routing group to the Asia routing group when all links between routing groups are functioning normally. the message travels through the least cost path through the North routing group. However, if the link between the North routing group and the Asia routing group is offline, the least cost path between the South routing group and the Asia routing group is through the Australia routing group.

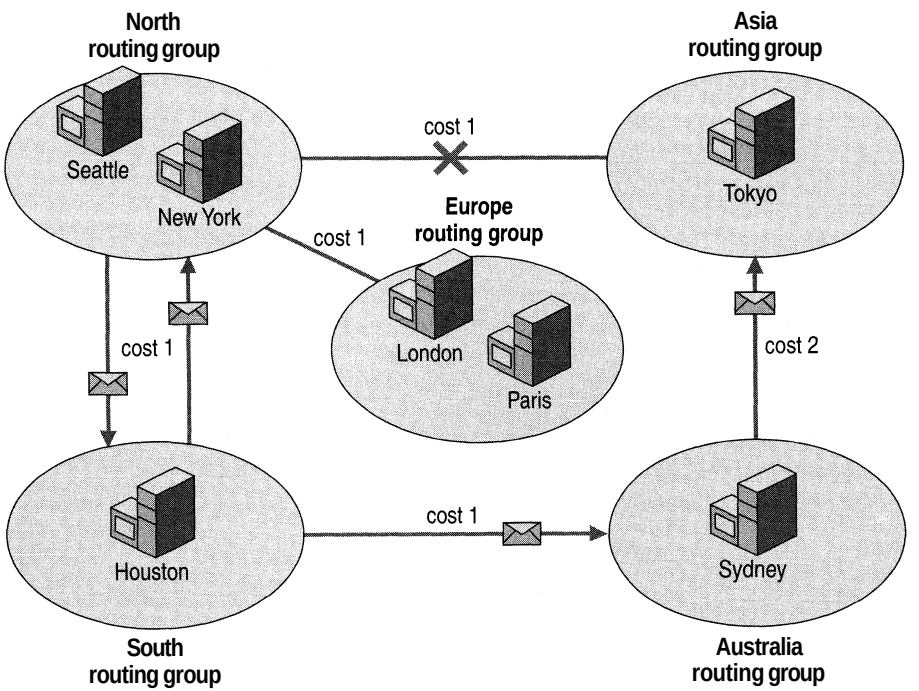


Figure 31.9 Rerouting when a connection fails between routing groups

The message reroute process works as follows:

1. The sender's server sends the message to the bridgehead server in the South routing group that then connects to the North routing group.
2. The bridgehead server in the North routing group attempts to open an SMTP connection to the bridgehead server in the Asia routing group. This connection fails.
3. If there are multiple destination bridgehead servers specified on the connector in the Asia routing group, the local bridgehead server in the North routing group attempts to open a connection between each of those servers until it finds a connection or no more servers available.
4. If the link between the North routing group and the Asia routing group is still not functioning after three connection attempts, it marks the connection as down and propagates the link state information.
5. The bridgehead server in the North routing group calculates an alternative route to the Asia routing group through the Australia routing group.
6. The message is rerouted from the North routing group to the Asia routing group through the Australia routing group.

Retries

The operating bridgehead server continues to try the connection until it is restored. You should consider the following conditions if you are tracking a message or managing a failed connection:

- If all routes to the destination routing group fail, the cost of the connection is set to infinite. The active bridgehead server checks on the status of the links three times at 60 second intervals and then retries according to the schedule set on the **Delivery** tab of the SMTP virtual server. If the failed connection becomes available again, it marks the connection as up in the link state table.
- Exchange holds messages destined for the failed routing group in the local message queues Exchange does not send messages through routing groups to a routing group that has no working connectors. If the expiration time-out specified on the **Delivery** tab of the SMTP virtual server passes 48 hours, the messages return to the originators as non-delivered.
- When Exchange sends a message to a connector or gateway for another system, such as MS Mail or cc:Mail, Exchange considers the message delivered when it reaches the connector, and rerouting does not occur, even if the other system connector cannot deliver the message.

- Exchange does not reroute a message routed to a connector with a remote initiated activation schedule unless the activation schedule changes before the message is delivered.
- Exchange stores a retry count on each message for each connector that it tries. When a message fails to connect to the remote system, Exchange increments the retry count and immediately reroutes the message. If Exchange cannot deliver the message after it tries all of the connectors or it reaches the maximum number of retries, it returns the message with a non-delivery report (NDR).

Multiple Public Folder Trees

With Exchange 2000 Server, you can use multiple public folder trees to ensure productivity, and you have many options to effectively manage them. Exchange 2000 Server provides administrative control and flexibility by supporting multiple public folder trees, also referred to as top-level hierarchies. For example, you can create a separate public folder tree to collaborate with external users and keep that content separate from the default public folder tree. Or, you can create an additional tree at a remote location for the users at that location to access data that is only relevant for them.

Each public folder tree stores its data in a single public folder store per server. You can replicate specific folders in the tree to every server in the organization that has a public folder store associated with that public folder tree.

For more information about public folders, see Exchange 2000 Help.

Client Support

When you install Exchange 2000 Server, it creates the default All Public Folders Tree. This tree is available to all MAPI, IMAP4, NNTP, and HTTP clients.

Additional public folder trees are only available to NNTP and Web clients, not to clients such as Outlook 2000 (unless viewed on a Web page hosted in Outlook 2000). You can use non-MAPI-accessible folders for collaboration with browsers and applications, such as Microsoft Office 2000, that can use HTTP to access the Web Storage System.

Support Considerations

You should consider the following when planning to support multiple public folder trees:

- Because Exchange creates the default public folder tree on every public folder server and it always replicates its list of folders, additional public folder trees only affect the servers on which they are configured. This means that you can create a set of departmental or local folders on only one server or a subset of servers. You need not replicate these additional public folders to every public folder server.
- You can use additional public folder trees to minimize the overall size of the default public folder tree, which simplifies navigation and reduces the cost of replicating the hierarchy of the default tree.

Latency Issues

The list of public folders in the global address list and Exchange System Manager is managed by different Exchange Server services than those that manage the replication of public folders. As a result, the following conditions might occur:

- When you administer the same public folder tree on two different servers, you do not see exactly the same list of folders. This means that changes to the public folder hierarchy, such as a new, renamed, or deleted folder, are not yet replicated among all servers.
- Public folders appear in the hierarchy with the client, but you cannot see them in the address book. If you cannot view a public folder in the address book, this means that the address list has not yet been regenerated because the folder is mail-enabled (or you might need to restart the client).
- The **Public Folder Properties** dialog box in Exchange System Manager (or the Exchange Folder snap-in) does not show the directory-specific pages. This means that the replication of the public folder hierarchy has taken place faster than the relocation of the newly-created directory information, including the address information.

Managing Public Folders

Managing public folders includes determining which public folders to mail-enable, and creating public folder replicas and referrals. The following sections describe these management tasks necessary to optimize your public folder hierarchy.

Public Folders in Active Directory

You can set up every public folder in a public folder store to appear as a mail recipient in Active Directory.

After you mail-enable a public folder the following conditions might occur:

- The System Attendant connects to Active Directory and creates an object for the public folder in a container such as Users. This container is specified on the **General** tab of the properties configuration of the public folder tree and applies to all public folders in the tree.
- A directory entry exists with the name: Folder Name + Global Unique Identifier. Users with access to Active Directory can use the e-mail address properties of the object to send email to the public folder.
- Additional tabs are available for the public folder in the Exchange System Manager, and the Active Directory Users and Computers snap-in. They are **E-mail addresses, Exchange general,** and **Exchange advanced.**
- You can configure the page to appear in the global address list for clients, such as Outlook.

  Note In Exchange Server 5.5, public folders are placed in the directory by default, but they do not display in the global address list. In an Exchange 2000 mixed-mode environment this does

not occur. New public folders are mail-enabled and configured as visible in the global address list. If you run Exchange 2000 in native mode, new public folders are not mail-enabled by default.

Public Folder Replicas

When you create a public folder, only one copy of the public folder exits within the organization. A public folder can exist in an organization either as a single copy or as multiple copies. Multiple copies are known as replicas. Using public folder replicas provides multiple, redundant information points in addition to load balancing for accessing data.

Creating a Public Folder Replica

A replica, which is copied from one server to another, is a separate instance of a public folder and its contents. For more information about creating public folder replicas, see Exchange 2000 Help.

Replicating System Folders

When designing your Exchange 2000 environment, you might need to create additional replicas of system folders to control your network traffic.

Table 31.2 describes common system folders that are created on a server when it is installed as the first server in an administrative group.

Table 31.2 Common system folders

Folder Name	Description
Reforms Registry	Storage for forms saved to the Organization Forms Library.
Events Root	Contains scripts for an Exchange Server 5.5–compatible event service.
Offline Address Book	Stores offline address books for clients to download.
Schedule+ Free Busy	Stores schedule information for clients to download.
Schema	Defines properties for objects kept in the public folder store.
StoreEvents	Contains Exchange 2000 event sink code for a specific server.
System Configuration	

Connecting to a Public Folder Replica

When a client attempts to access public folder data, the client must be able to connect to a server that contains a replica of the data. The client attempts to connect to any replica to present the requested data to the user.

To maximize efficiency, the client attempts a connection to servers in the following order:

1. The default public folder store for the clients. The default public folder store is determined by the configuration of the mailbox store containing the user's mailbox. If the default public folder store is not available, the client receives a list of servers that contain the replica.
2. A server to which the client has an existing connection.
3. Each server within the same server routing group as the public folder server routing group for the client.
4. If the client cannot connect, Web Storage System instructs the client to attempt a connection to other routing groups in the order of the routing group connection values.
5. Because the connections to routing groups have the same cost, the servers containing the replicas are pooled together and selected at random as if they were in the same routing group.

Public Folder Referrals

Public folder referrals allow you to route information and requests to specific folders.

You can enable public folder referrals to servers in another routing group by implementing and configuring a routing group connector between the two routing groups. The Routing Group connector is one directional; it requires that two instances be configured for bi-directional traffic. You can configure public folder referrals for the routing group going in each direction.

Public folder referrals between routing groups are transitive and allow all referrals over the connection when enabled. For more information about creating public folder referrals, see Exchange 2000 Help.

CHAPTER 32

Real-Time Collaboration

Traditional e-mail systems do not allow transfer of dynamic data and information on demand. Real-time collaboration services such as Microsoft Exchange Instant Messaging Service, Microsoft Exchange Chat Service, and Microsoft Exchange 2000 Conferencing Server provide the immediacy of the telephone, with the features of e-mail.

In This Chapter

Instant Messaging

Chat Service

Exchange Conferencing Server

Instant Messaging

Instant Messaging allows users to see presence information (such as online, offline, out-of-the-office, busy) for other users and provides the ability to send instantaneous communications. For example, members of a team that are collaborating with one another might need the convenience of information delivered instantaneously.

Instant Messaging Components

The Instant Messaging architecture includes components located at both the client and the server. The Instant Messaging server itself runs as part of the Microsoft Internet Information Services (IIS) process (Inetinfo.exe) and is implemented as an Internet Server Application Programming Interface (ISAPI) extension. Figure 32.1 illustrates the architecture of Instant Messaging.

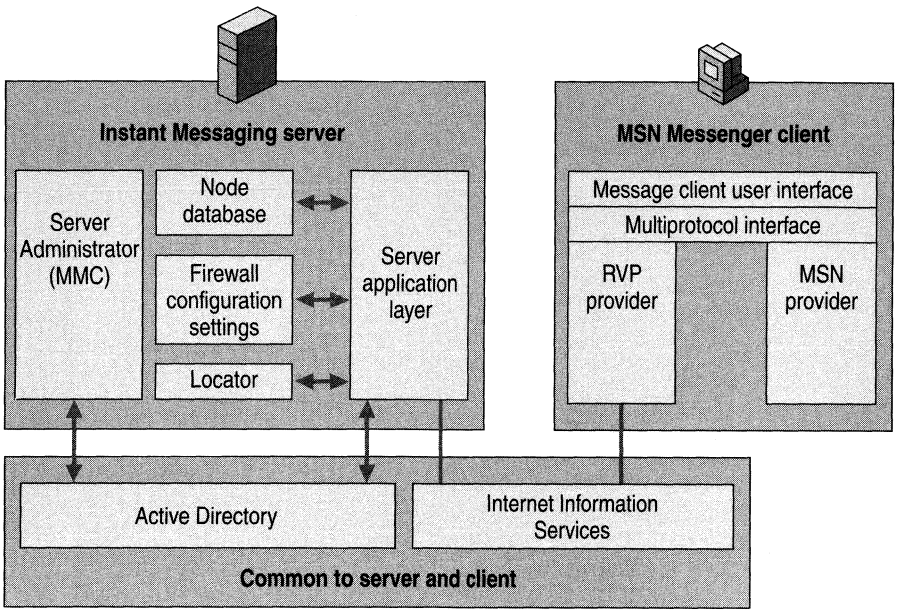


Figure 32.1 Instant Messaging architecture

The following sections describe the components involved in Instant Messaging on the server and on the client, as along with the Microsoft Windows 2000 Server features and components used by Instant Messaging.

For more information about Instant Messaging, see Exchange 2000 Server Help.

Instant Messaging Server Components

The Exchange Instant Messaging Service runs as part of the Windows 2000 Internet Information Services (IIS) process (Inetinfo.exe) and it is implemented as an ISAPI extension. Figure 32.2 shows the Instant Messaging virtual server.

> **Note** The Instant Messaging object located under the **Protocols** container in System Manager is a representation of an Instant Messaging virtual server. Each Instant Messaging virtual server requires the presence of one IIS virtual server.

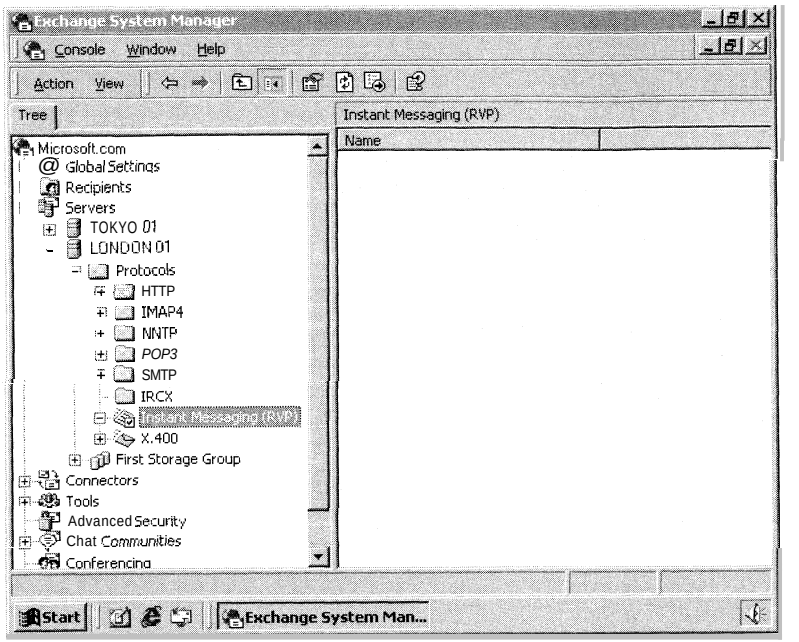


Figure 32.2 Instant Messaging virtual sewer window

Sewer Application Layer

The Server Application Layer performs most of the Instant Messaging processing and communicates with other server-side Instant Messaging components and with Active Directory.

RVP

All Instant Messaging communication between clients and servers uses a protocol called RVP, which is an extended subset of Hypertext Transfer Protocol (HTTP) with Extensible Markup Language (XML) features.

Node Database

The node database is an instance of the Extensible Storage Engine (ESE) that retains subscription information for Instant Messaging clients on the local server. This information, commonly known as a subscription list, is a list of clients who have subscribed to a user on the local server.

For example. Jim's local server is home server 1; Alice decides to subscribe to Jim's information—that is, she wants to be able to watch his online status. Jim's subscription list on the node database of home server 1 maintains a list of subscribers to Jim's information. When Alice subscribes to Jim's information, her Instant Messaging user address is added to Jim's subscription list. When Jim's status changes, home server 1 sends this status change notification to Alice and she is instantly notified of Jim's status change.

The node database makes it possible for all users who subscribe to a user's information to receive up-to-date status information without congesting the network with the traffic created by the user's home server continually checking with all its subscribed users to maintain current online status.

Locator

You use the locator to dispatch notifications to the correct home server when messages go through the Instant Messaging router.

Redirection

The Instant Messaging server also redirects clients within the organization to home servers, and it performs proxy operations for clients outside a network firewall.

Instant Messaging Routers

You can deploy Exchange 2000 Server in a front-end and back-end architecture where you can split the protocols and Microsoft Web Storage System onto different servers. Instant Messaging servers operate in a similar way.

You can use a bank of lnstant Messaging routers as the primary point of contact to redirect clients to the home server for the user. The lnstant Messaging router configuration allows one or more front-end routers to provide a unified namespace.

Firewall Topology Settings

The firewall settings in System Manager retain information about each Instant Messaging server and whether a client is inside or outside the firewall in relationship to the local server. The **Firewall Topology** tab contains data that tells whether a particular source Internet Protocol (IP) address (or client) can connect to a particular destination IP address (or server), and whether a proxy server is required.

> **Note** The firewall topology setting that determines whether a proxy server is required is active only on the server side of the connection.

An Instant Messaging router can perform three possible actions to manage a transaction:

- Relay the transaction. The server itself connects to the destination and relays the results back to the sender.
- Redirect the transaction. The server instructs the sender to connect directly to another destination that it specifies.
- Reject the transaction. If, based on the network topology and the source and destination IP addresses of the request, the server cannot serve process the transaction locally, it rejects the transaction.

Figure 32.3 provides a step-by-step illustration of what happens when the Instant Messaging router relays a transaction.

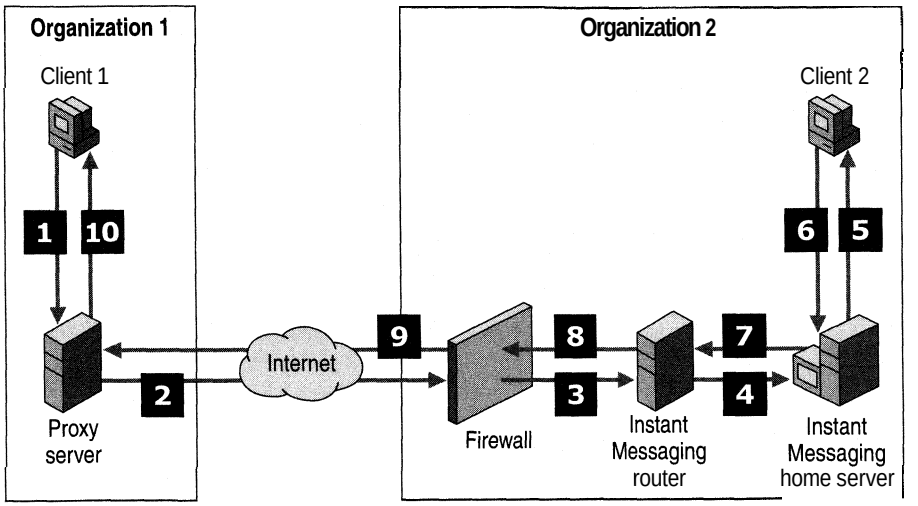


Figure 32.3 Instant Messaging router relays response to a request

In this example, Client 1, who belongs to Organization 1, sends an Instant Message request to Client 2, who belongs to Organization 2. The following steps describe the action taken:

1. Client 1's client application determines that the Instant Messaging recipient is outside Organization 1's domain and sends the message to a proxy server.
2. The proxy server sends the Instant Message over the Internet to Organization 2, where it encounters Organization 2's firewall.
3. The Instant Message transfers across the firewall to an Instant Messaging router in Organization 2.
4. The Instant Messaging router in Organization 2 consults its routing table and finds that Client 1 cannot connect directly to Organization 2's home server because the server is not exposed to the Internet. Instead, the Instant Messaging router acts as a proxy for Organization 2's home server and gateways the message to Organization 2's home server.
5. Client 2's home server directs the Instant Message to Client 2.
6. Client 2 responds to the Instant Message.
7. The client response travels to the Instant Messaging router.
8. The client response travels through the firewall.
9. The client response travels across the Internet to the proxy server in Organization 1
10. The client response travels to Client 1.

Instant Messaging uses RVP, which is an extended subset of HTTP. Each transaction in HTTP includes a message and response code pair. Thus, there is always a response to an instant message, and the connection exists until that response code returns. For example, in Figure 32.3 five

connections are opened: client to proxy, proxy to firewall, firewall to router, router to home server, and home server to client. Before Client 1 sends its response back, all five connections are open. When Client 1's home server receives the response code back from the client, it terminates the home server to client connection. Once the router receives the response code from the home server, the router to home server connection terminates, and so on back to the originating client to proxy connection, which is the last to terminate.

Figure 32.4 provides a step-by-step illustration of what happens when the Instant Messaging router redirects a transaction in response to a request.

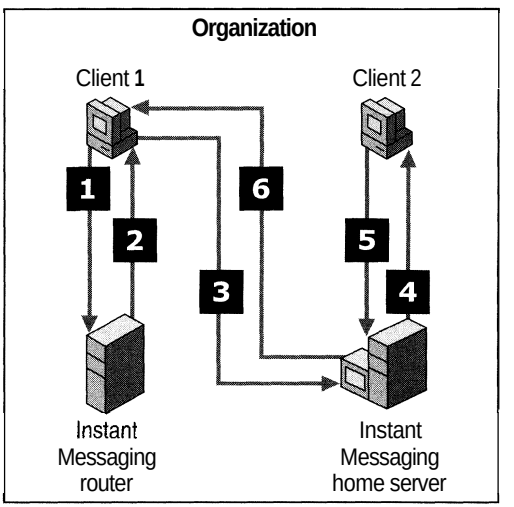


Figure 32.4 instant Messaging router redirects response to a request

In this example, Client 1, who belongs to the same organization as Client 2, sends Client 2 an Instant Message request. The following steps describe the action taken:

1. Client 1's Instant Message goes to the Instant Messaging router.
2. The Instant Messaging router consults its table and finds that Client 2 is in the same organization as Client 1. The Instant Messaging router redirects Client 1 to the home server.
3. Client 1 disconnects from the Instant Messaging router and then creates a connection to Client 2's home server.
4. The home server sends the message to Client 2.
5. Client 2 sends a response to the home server.
6. The home server retains Client 2's URL information from the initial request, so it sends the response directly to Client 2 without consulting the router.

Instant Messaging Client Components

You must use an Instant Messaging client to log on to an Instant Messaging server and communicate with other Instant Messaging users. An Instant Messaging client includes the following components:

- **Instant Messaging client user interface** Users can use the Instant Messaging client to log on to an Instant Messaging server, configure security, and configure support for other messaging providers.
- **Providers** The Microsoft Network (MSN) provider allows the Exchange Instant Messaging client to comniunicate with contacts on MSN.
- **RVP support** The Instant Messaging client uses the RVP protocol to communicate with Exchange 2000 Instant Messaging servers.

Client Operations

The client is an important part of the Instant Messaging system. The user logs on the client, and the Instant Messaging client caches the user URL. The Instant Messaging client maintains the user's status and contacts on the server.

Logon

When users log on to Instant Messaging, they provide their Instant Messaging user ID in the form of *username@imdomain* for authentication on their home server. If the Instant Messaging service and the user accounts that use Instant Messaging are in different forests, the user must provide Windows 2000 user names and passwords.

A number of operations occur to authenticate users when they log on. Figure 32.5 shows the client logon process.

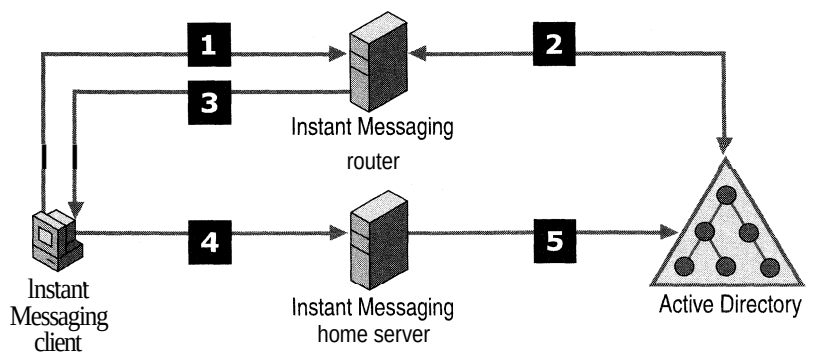


Figure 32.5 Client logon process

The following steps illustrate the client logon process:

1. The user logs on to the client with *username@imdomain* and connects to the Instant Messaging router.
2. The Instant Messaging router queries Active Directory for the user's Instant Messaging home server.
3. The Instant Messaging router returns the home server URL to the client.
4. The Instant Messaging client uses the home server URL to connect to the Instant Messaging home server.
5. The Instant Messaging home server queries and validates the user's Active Directory user name and password.

   Note Microsoft Windows NT users must place the network basic input/output system (NetBIOS) name of the Instant Messaging server in the Domain field of the Instant Messaging client if they are located in a different forest. (This is the same as the Windows 2000 Domain field entry.)

Client Caching

Once a client locates another user, it does not go through the router to access that user again. The client caches the URL and continues to use it until the application shuts down and clears the cache.

Contact Subscriptions

Instant Messaging users can add contacts to their client. After a contact is added, the Instant Messaging client receives status information from the contact's home server. This is known as a *subscription.* A contact's subscriptions are stored in the node database on the user's home server, and users maintain the list of other users they subscribe to on their local computer in the client registry under the following key:

```
HKEY_CURRENT_USER\Software\Microsoft\Exchange\messenger\profiles\
http://imdomain/instmsg/aliases/alias\Contacts
```

Instant Messaging Addressing

All Instant Messaging users are identified by unique URLs. Each user has two URLs: an Instant Messaging URL and an Instant Messaging home server URL. When logging on to Instant Messaging or adding contacts, users can use a simplified address based on a Simple Mail Transfer Protocol (SMTP) address called the Instant Messaging user address.

To access the Instant Messaging URL, Instant Messaging home server URL, or Instant Messaging user address, you use the Active Directory Users and Computers console in Microsoft Management Console (MMC).

To display a user's Instant Messaging property sheet

1. In Active Directory Users and Computers, expand the server on which the desired user's Instant Messaging user address resides.
2. Click the **Users** folder, and then right-click the user.
3. Click **Properties,** and then click the **Exchange Features** tab.
4. Click **Instant Messaging** and then click **Properties.**

This displays the **Instant Messaging** properties shown in Figure 32.6.

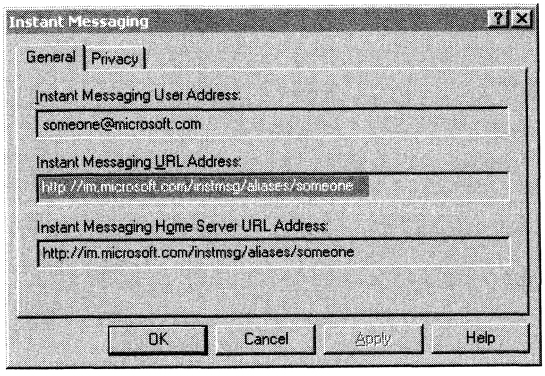


Figure 32.6 Instant Messaging properties

Instant Messaging URL

The Instant Messaging URL points to the Instant Messaging router for the Instant Messaging domain. For example, the user Bob has the following domain URL:

```
http://%routername%/instmsg/aliases/bob
```

You can consider this URL public because those outside the Instant Messaging domain can see it.

Instant Messaging Home Server URL

The Instant Messaging home server URL points to the Instant Messaging home server. For example, the user Bob at Microsoft who is an Instant Messaging user on the Vancouver server in an Exchange organization with an Instant Messaging domain equal to irn.rnicrosoft.com has the following home server URL:

```
http://imhomeserver/instmsg/local/im.microso£t.com/instmsg/aliases/bob
```

You can consider this URL private because only those inside the Instant Messaging domain can see it.

> **Note** The Instant Messaging home server URL allows a user to retain the same Instant Messaging user address, even if the home server changes.

Instant Messaging User Address

A more convenient format than URLs to identify Instant Messaging users is an Instant Messaging user address. An Instant Messaging user address is formatted the same way as a standard SMTP e-mail address. To be easily accessible to other users, an Instant Messaging user address should use the following naming convention:

```
username@imdomain
```

For example, the user Bob at Microsoft Corporation has the following Instant Messaging user address:

```
bob@im.microsoft.com.
```

> **Note** You can configure the lnstant Messaging domain to be the same as the SMTP domain by using Domain Name System server resource records (DNS SRV). In this example, the resulting address is bob@microsoft.com. For more information about using DNS SRV records, see "Configuring DNS" later in this chapter.

Windows 2000 Dependencies

Exchange 2000 Instant Messaging depends on Windows 2000 for several important functions that are provided by Active Directory and IIS version 5.0.

Active Directory

You can create Instant Messaging user accounts by using the Active Directory Users and Computers console. It is not necessary to have a separate user account list for Instant Messaging. All users in Active Directory can use Instant Messaging.

Active Directory also provides security when connecting to Instant Messaging servers by requiring users to provide their Windows 2000 user ID and password for authentication before logging on to their Instant Messaging server.

Internet Information Server 5.0

Because Instant Messaging runs as an Internet Server API (ISAPI) extension, it is depends on IIS 5.0 for Internet functionality.

Configuring Instant

Instant Messaging Service installs with Exchange 2000. During installation, the Active Directory schema is updated with some new classes and attributes to support the Instant Messaging infrastructure.

Table 32.1 lists the new classes and attributes added to Active Directory when Instant Messaging installs.

Table 32.1 Classes and attributes added to Active Directory

Class	Attribute	Data Type	Description
User/*name*(auxiliary class ms-Exch-IM-Recipient)	msExchIMACL	<binary blob>	Global/Access Control List
	msExchIMAddress	<string>	Global/Instant Messaging user address (e-mail-like form)
	MsExchIMMetaPhysicalURL	<string>	Global/Public URL of user
	MsExchPhysicalURL	<string>	Global/Physical URL of user (computer based)
	MsExchIMVirtualServer	<string>	Global/Home server for user
RVP (ms-Exch-Protocol-Cfg-IM-Container)	MsExchIMDBLogPath	<string>	Path of DB log file
	MsExchIMDBPath	<string>	Path of DB file
RVP/1 (ms-Exch-Protocol-Cfg-VIM-Virtual-Server)	MsExchIMServerHostsUsers	TRUE/FALSE	Determines if server can host users
	MsExchIMServerIISID	<digit>	Identifier of IIS hosting server
	MsExchIMServerName	<string>	Name of server
	MsExchIMHostName	<string>	DNS Host name of virtual server
Instant Messaging (ms-Exch-IM-Global-Settings-Container)	None	None	None
Instant Messaging/1 (ms-Exch-IM-Firewall)	MsExchIMFirewallType	<digit>	Type of firewall
	MsExchIMIPRange	<string>	Semicolon-separated range of IP address
	MsExchIMProxy	<string>	DNS name of proxy

To configure Instant Messaging, you need to configure DNS and the server components.

Configuring DNS

When users enter their Instant Messaging user address into the client logon screen, the client transforms the string into a URL. So, user@im.microsoft.combecomes http://im.microsoft.com/instmsg/aliases/user. The client uses DNS to locate the physical server for the operation.

You should base the Instant Messaging domain of a user's Instant Messaging user address on the SMTP hostname of the user's e-mail address. For example, a user whose e-mail address is user@vancouver.rnicrosoft.comshould have an Instant Messaging domain of im.vancouver.microsoft.com.

Using the im prefix to denote an Instant Messaging domain allows external users to easily determine an Instant Messaging user address from an e-mail address. For example, users whose e-mail addresses end with @rnicrosoft.comshould have the Instant Messaging domain @im.microsoft.com. Similarly, users whose e-mail addresses end with @jp.microsoft.com should have the Instant Messaging domain @im.jp.microsoft.com.

Unified Namespace

You can simplify Instant Messaging users' addresses by using DNS SRV records to create a unified namespace, so that a user's Instant Messaging user address is the same as the SMTP address.

There are three possible scenarios for the use of DNS SRV records:

- Do not use DNS SRV records.
- Use DNS SRV records but do not expose them externally. This requires that users have two addresses: one with a DNS SRV record for use internally and one without a DNS SRV record for use externally.
- Use DNS SRV records internally and externally.

Figure 32.7 shows the difference between a separate namespace and a unified namespace.

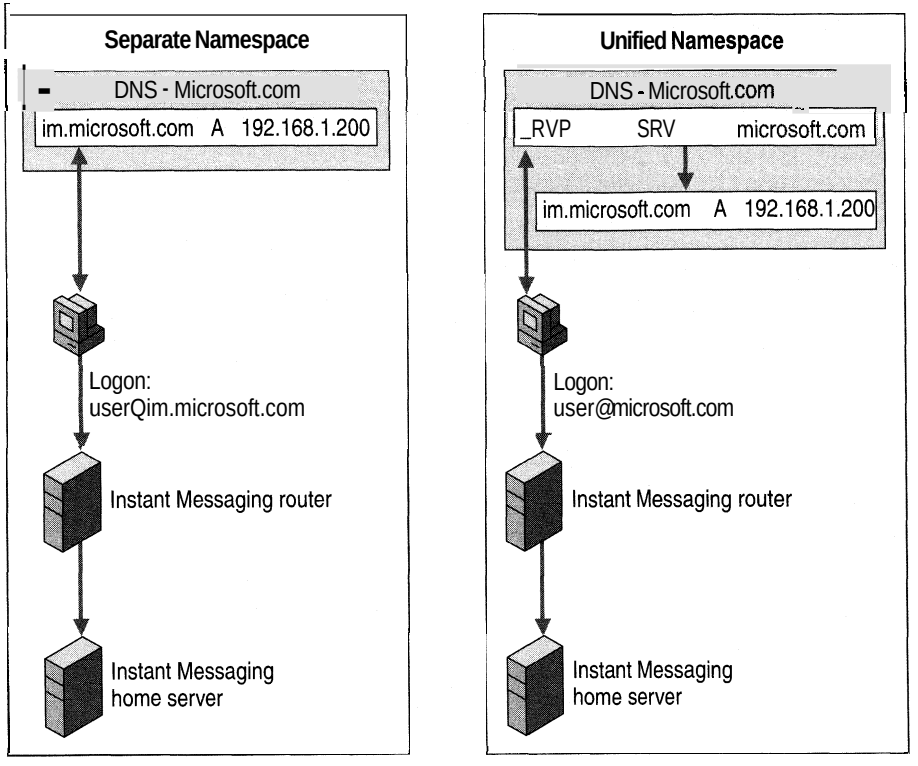


Figure 32.7 Comparison of separate and unified namespaces

DNS SRV Records

The DNS SRV record allows Instant Messaging to query DNS for the host name for a particular service. For example, if Joe in the Vancouver office of Microsoft Corporation wants to exchange instant messages with Suzan in Japan, whose e-mail address is suzan@jp.microsoft.com, Joe adds the SMTP address suzan@jp.microsoft.comto his Instant Messaging contact list. Joe's client performs a DNS SRV lookup at the jp.microsoft.com DNS zone for the RVP service, and learns that the Instant Messaging router for this zone is im.jp.microsoft.com. Based on this information, Joe's client constructs the Instant Messaging URL http://im.jp.microsoft.com/instmsg/aliases/suzan. This is how the SRV record allows Joe to communicate with Suzan by simply using Suzan's e-mail address.

It is recommended that you consider the following guidelines when planning your Instant Messaging network:

- Create home server and Instant Messaging routers on different physical servers.
- If you plan to connect your Instant Messaging network to the Internet, install all home servers and Instant Messaging routers on computers protected by a firewall.
- Do not allow direct Internet connections to home servers. Instead, use one or more Instant Messaging routers as front-end servers to handle Internet traffic. This protects user accounts and data on the home servers and avoids exposing host server names to the Internet.
- Direct all outgoing presence information requests and instant messages through an HTTP proxy server.
- Direct all incoming presence information requests and instant messages though TCP port 80 on a reverse proxy server. This ensures the privacy of your Instant Messaging router IP addresses.
- Use a consistent and intuitive naming convention for your Instant Messaging servers, domains, and addresses.
- Use DNS SRV records so that you have a unified namespace with your SMTP address

  Note Microsoft Windows 98 or earlier and Windows NT clients must have Active Directory client installed before they can search by first name or last name. Active Directory client is also required if they want to perform DNS SRV record lookups, which mask the Instant Messaging server's computer name. Users of these operating systems also need Windows Sockets (WinSock) 2.0 and Microsoft Internet Explorer 5 to use Instant Messaging.

  You can find Active Directory client for Windows 98 or earlier on the Windows 2000 CD, and you can find Active Directory client for Windows NT 4.0 on Windows NT 4.0 Service Pack 7.

Configuring Server Components

You configure Instant Messaging virtual servers and the firewall topology by using the Exchange System Manager console in MMC. You configure the firewall topology with the Instant Messaging settings under **Global Settings** for the Exchange 2000 organization in Exchange System Manager. The following sections describe how to configure the Instant Messaging virtual servers and routers.

Short Names and Fully Qualified Domain Names

A short name is usually a single word, such as Tokyo-01. A fully qualified domain name (FQDN) fully describes the location of a virtual server in the organization and includes the domain name — for example, Tokyo-01.jp.microsoft.com. You use either short names or FQDNs when setting up Instant Messaging virtual servers.

It is recommended that you name all routers by using FQDNs. It is also recommended that you name home servers by using short names unless Instant Messaging is installed in a forest with multiple trees. If you install Instant Messaging in a forest with multiple trees, name home servers by using FQDNs. One of the benefits of using short names is that no proxy server is necessary.

Note If you use the Instant Messaging Virtual Server Wizard, the default virtual server name is a short name.

Instant Messaging Virtual Servers

Tnstant Messaging virtual servers are configured either as routers or as home servers. The distinction between a router and a home server is that only home servers can host users; routers cannot. This is because you might need to set up your network configuration with multiple routers with the same virtual server name. DNS cannot locate an Instant Messaging user with accounts on multiple virtual servers.

Note It is recommended that you do not make home servers visible to the Internet. Only routers should be visible to the Internet. It is also recommended that you create a DNS SRV record for routers so that your Instant Messaging user address looks like your e-mail address. For more information about making addresses identical, see "Configuring DNS" earlier in this chapter.

Firewall Topology

To configure the firewall topology, open System Manager open **Global Settings,** and then right-click **Instant Messaging Settings.**

You can configure your firewall to protect a range of IP addresses. Thus, IP addresses within a certain range are considered inside the firewall. You can also configure the address of the HTTP proxy server that you use for outbound connections.

Instant Messaging Server Limits

The number of servers required by an organization grows in accordance with the size of the user population. Depending on your hardware, a single Instant Messaging home server can handle up to 10,000 users that are concurrently online, and a single Instant Messaging router can handle up to 50,000 users that are concurrently online. A user is considered online when the client is actively running on the desktop and a user is logged on to the Instant Messaging system. A typical corporate user is logged on to the system during the workday. Some users, such as users who leave their clients running all the time, are logged on to the Instant Messaging system continuously. The client is marked as idle when not in use.

In typical corporate settings, approximately 80 percent of all users are concurrently online at peak times. In Internet service provider (ISP) settings, typically only 5 to 10 percent of all users are concurrently online.

Therefore, a company with 30,000 users who are online most of the work day needs three home servers, because almost all 30,000 users might log on at peak times. An Internet service provider (ISP) with 1,000,000 users and a 5-percent peak online rate needs five home servers, because only 50,000 users are ever concurrently online.

Instant Messaging Deployment Scenarios

You can use a variety of strategies to deploy Instant Messaging. The following deployment models can help you determine the best strategy for your needs.

Small Business Deployment

If you have fewer than 10,000 users who are online simultaneously and if equipment and IT personnel are limited, perform a small business deployment. This deployment is not complex, does not span multiple domains, and is not expected to span multiple domains in the future.

In this deployment, you use only one Instant Messaging home server and there is no need for Instant Messaging routers. Firewalls are optional.

Figure 32.8 illustrates a possible Instant Messaging deployment configuration for a small business.

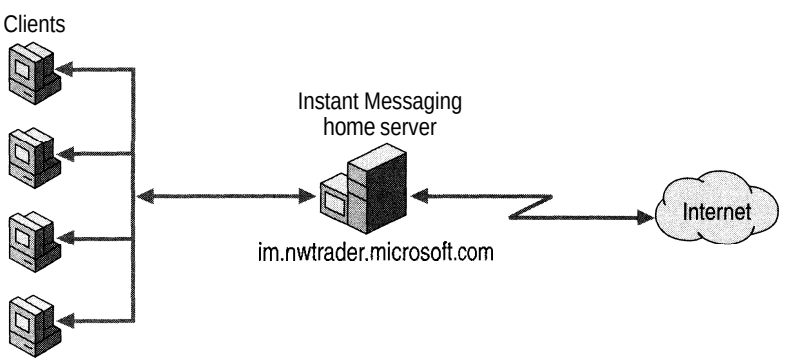


Figure 32.8 Small business deployment

Standard Deployment

The standard deployment is effective in an environment exceeding 10,000 users who are online simultaneously. It provides a unified namespace, hides complete deployment details from users, and can span multiple domains. It can be a geographically diverse deployment, and it is visible to the Internet. The configuration can use multiple Instant Messaging routers; it also uses firewalls for inbound security and HTTP proxies for outbound security.

Figure 32.9 illustrates a standard Instant Messaging deployment configuration.

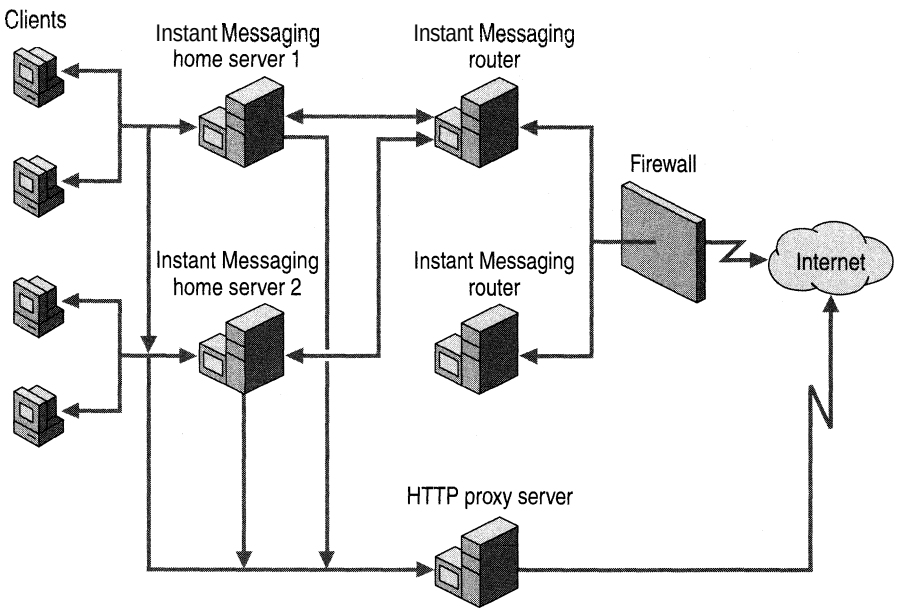


Figure 32.9 Standard Instant Messaging deployment

Enterprise Deployment

The enterprise deployment is a variation of the standard deployment, except that a perimeter *network* exists. A perimeter network (also called a demilitarized zone [DMZ]) is a computer host or small network inserted as a neutral zone between a company's private network and the outside public network.

In this configuration, routers in the perimeter network have direct Internet connectivity. This deployment provides higher security than the standard deployment.

Figure 32.10 illustrates an enterprise Instant Messaging deployment configuration.

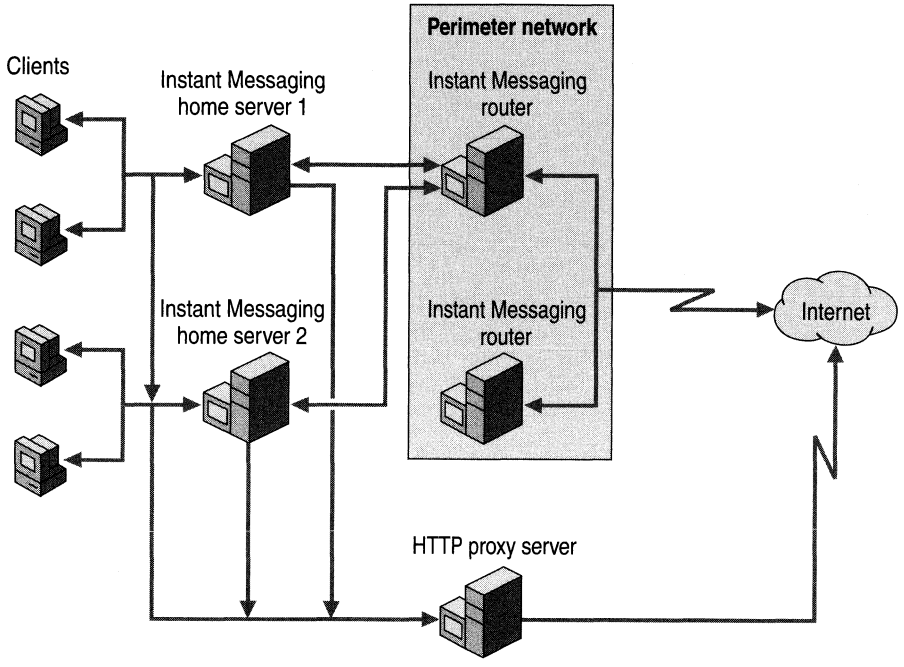


Figure 32.10 Enterprise deployment

Multiple E-mail Domain Deployment

The multiple e-mail domain deployment is a variation of the standard deployment. With this deployment, you can allow multiple commonly available e-mail addresses to be used externally (such as user-a@na.rnicrosoft.com and user-b@jp.microsoft.com) and you can provide IT service for multiple companies.

In this deployment, different Instant Messaging routers answer to different Instant Messaging domains, and an Instant Messaging home server can service multiple domains.

Figure 32.11 illustrates an Instant Messaging deployment configuration for multiple e-mail domains.

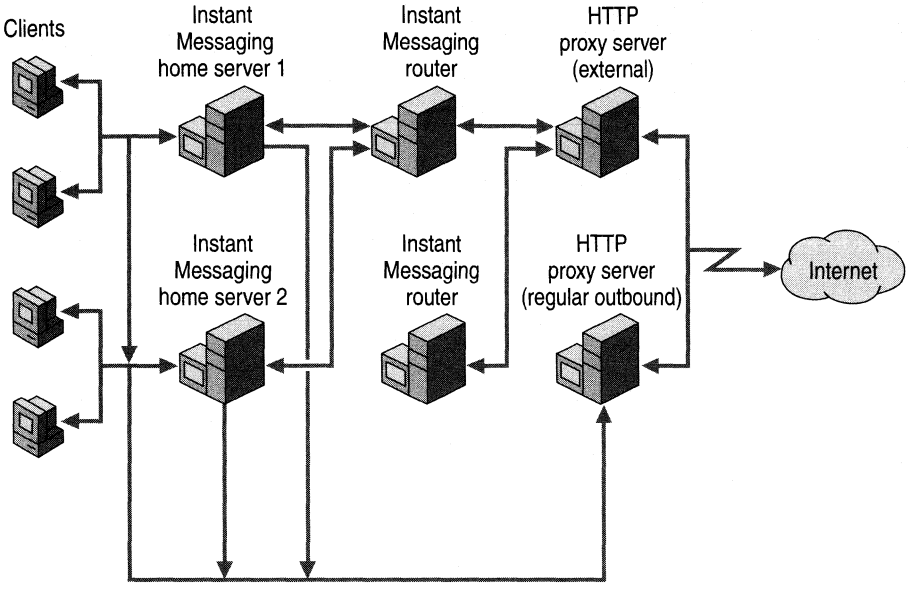


Figure 32.11 Multiple e-mail domain deployment

ISP Deployment

In the ISP deployment configuration (not in hosting configurations), several million users can reside on the Instant Messaging home servers at one time because there is a low ratio of expected concurrent online users.

Multiple Instant Messaging home servers and multiple Instant Messaging routers are required, but no firewalls exist between Instant Messaging routers and the Internet.

Figure 32.12 illustrates an ISP (non-hosting) Instant Messaging deployment configuration.

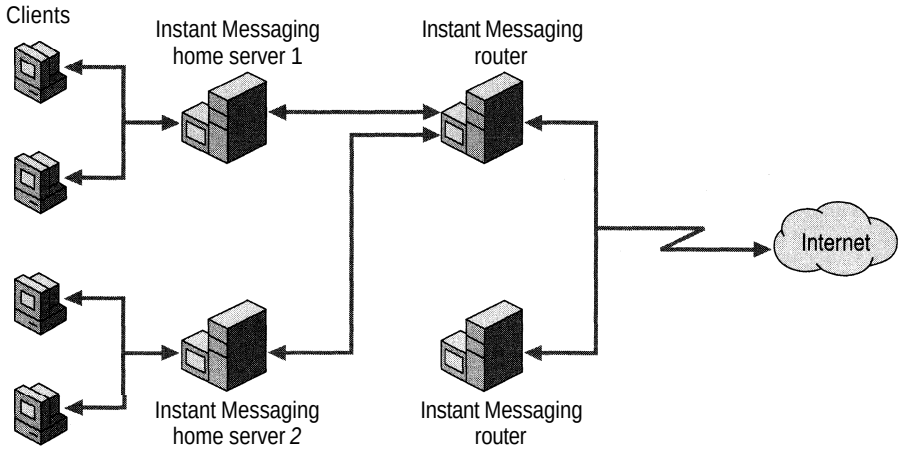


Figure 32.12 ISP deployment

Chat Service

Unlike Instant Messaging, which is designed primarily for one-to-one communications, Chat Service allows users to join a chat room and communicate with other users in a forum. There are additional features in Chat Service that allow users to "whisper" messages to one another, although this is normally the exception rather than the rule.

You can use Chat Service to hold organizational meetings where users are spread over long distances. If you don't want all the participants asking questions at the same time, you can have a structured question-and-answer session. You can also use Chat Service for brainstorming sessions, where users can submit their thoughts without being explicitly invited to a conference.

Server Configuration

The Exchange 2000 Chat Service allows you to install a fully functional chat community within minutes. You can install Chat Service on Exchange 2000 Server by using the standard installation program. You only need to define some channels (or *rooms* as they are known in the client) to which the users can connect.

You can configure Chat Service with some advanced options, such as bans that restrict certain clients from accessing a channel. You can also configure each channel with a Platform for Internet Content Selection (PICS) rating so that users can have some forewarning regarding the content and language in the room.

Scalability

Each channel can accommodate up to 5,000 users, although you can adjust this setting as required. Due to the nature of the Chat Service protocol, an entire channel must be located on a single chat server. However, Exchange 2000 Chat Service is highly scalable and you can deploy it in large companies or Internet service providers.

Client Connections

Because the chat protocol is an open industry standard, many types of clients from different vendors can connect to an Exchange 2000 Chat Service server, such as Microsoft Comic Chat version 2.0. Users generally create their own local profiles, defining their real names, aliases, and other information. Users must also specify the name of the chat server to which they want to connect. Internally, a client connects to port 6667 on the chat server by default. From here, a client can list the current rooms or create new rooms with the necessary access permissions.

Exchange Conferencing Server

Microsoft Exchange Conferencing Server allows users to host virtual meetings and conferences through the use of audio, video, shared whiteboard, direct file transfer, and chat. This makes use of all multimedia facilities available on the client. Of all the real-time components in Exchange 2000, this is by far the most advanced.

To allow full-function conferencing, several server-side components in Exchange 2000 perform conference management and session coordination. You can schedule online meetings and reserve virtual resources through integration with Microsoft Outlook 2000. The client uses the T.120 protocol, which is integrated into products such as Microsoft NetMeeting, to join the conference and communicate with the conferencing server.

Exchange Conferencing Server consists of a resource reservation agent and a conference controller known as Conference Management Service. Together, these components provide users with the ability to schedule an online meeting against a conference resource.

Exchange Conferencing Server controls the Data Conferencing Provider. The Data Conferencing Provider provides users with server-based T.120 data conferencing, which is compatible with application sharing, white boarding, and chat services provided by client programs such as Microsoft NetMeeting.

Data Conferencing Provider also depends on Windows 2000 Server to provide support for IP Multicast Conferencing, security, and quality of service.

For more information about Exchange Conferencing Server, see *Exchange 2000 Conferencing Server Concepts and Planning* and Exchange 2000 Conferencing Server online documentation.

CHAPTER 33

Troubleshooting

Microsoft Exchange 2000 Server provides a variety of diagnostic tools that you can use to troubleshoot problems in your organization. This chapter addresses some of the problems you could encounter while using Microsoft Exchange 2000 Server.

In This Chapter

Installation and Setup Problems

Problems With Active Directory and Active Directory Connector

Web Storage System Problems

Connectivity Problems

Windows 2000 Server Tools Problems

Backup and Restore Problems

Performance Problems

Installation and Setup Problems

When you install and set up Exchange 2000, you could experience command failures or error messages.

To troubleshoot Exchange 2000 installations

1. If the ForestPrep command **(setup Iforestprep)** doesn't work, make sure the account has the following permissions:
   - Schema Admins
   - Enterprise Admins
   - Local Administrator rights to the server
2. If the DomainPrep command (**setup/domainprep**) doesn't work, make sure the account has the following permissions:
   - Domain Admins
   - Local Administrator rights to the server

If Setup fails while installing Exchange 2000, in most cases you should reinstall Exchange 2000. Network problems can cause setup failures. In these cases, first verify that the network connection is working correctly, and then complete the following steps to reinstall Exchange components.

To reinstall Exchange components

1. Turn off all Exchange services.
2. Restart the server.
3. Run the Setup program, and reinstall any components that Setup did not install successfully during the initial installation.

Setup Fails: "Setup Was Unable to Bind to the Exchange Server" Error Message

You attempt to join a computer running Microsoft Exchange 2000 Server to an existing Microsoft Exchange Server 5.5 site, and you receive the following error message:

```
Setup was unable to bind to the Exchange Server
```

This occurs when the account that you use to start Exchange 2000 Setup does not have the correct permissions for the target server's Exchange Server 5.5 directory.

Solution

Give this account Administrator permissions at the Site and Configuration levels of your Exchange 5.5 server.

Setup Fails When You Install Exchange in a Child Domain

You install Exchange 2000 Server into a child domain, and Setup stops working the first time with the following error message:

```
The Directory Service is busy
```

This problem occurs because of the way the named pipe transport functions. The server side creates an instance of the pipe for use by clients. The first client to attempt a connection is associated with that instance. To allow another client to connect, the server must create another instance of the named pipe. If another client attempts to connect before the new instance is created, the server appears (to the second client) to not accept a connection.

Solution

Start Exchange 2000 Setup again, then install the components that did not install the first time. After you complete these steps, Exchange 2000 Setup will finish successfully.

Setup Fails: "Error 0xC103798A" Error Message Displays

You install Exchange 2000 Server, and you receive the following error message:

```
Setup failed while Configuring registry entries for Exchange System
Management Snap-ins (error 0xC103798A: An internal component has failed.)
```

The command **regsvr32.exe -s maqadmin.dll** fails, returning error code 3 (The system cannot find the path specified).

If any component of Exchange Server 5.5 is previously installed on this server (and does not completely uninstall), that version of the Exchmem.dll file causes this conflict.

Solution

Before you run Exchange 2000 Setup, rename the existing Exchmem.dll to Exchmem.old. Exchange 2000 copies the appropriate Exchmem.dll file during Setup.

Setup Fails While Trying to Join Existing Exchange 5.5 Site

You attempt to install Exchange 2000 Server to join to an existing site that has a large number of servers or sites, and Setup fails with the following error message:

```
Idispatch not found for [Microsoft Exchange conferencing mmc snap-in]
```

Solution

Before you install Exchange 2000, complete the following steps on the Exchange *5.5* computer to which you are joining:

To install Exchange 2000 on an Exchange 5.5 computer

1. Run the Microsoft Exchange Server Administrator program.
2. Click to expand **Organization,** click to expand **Site,** click to expand **Configuration,** and then click to expand **Protocols.**
3. Double-click LDAP **protocol** to open the properties.
4. Click the **Search** tab and modify the value for **Maximum number of search results returned** to a value that is equal to or greater than the number of servers in the organization. If the LDAP protocol object at the server level is not set to use the site defaults, you should change the value at the server level.

Setup Fails While Joining an Exchange Server 5.5 Site on a Windows 2000 Domain Controller

You run Exchange 2000 Server Setup on a member server on a Microsoft Windows 2000 Server domain, and Setup does not work when the Windows 2000 Server domain controller is upgraded from Microsoft Windows NT 4.0 and has Exchange Server *5.5* installed. Setup doesn't work when you try to join the existing Exchange Server *5.5* site.

This problem occurs because Exchange Server *5.5* and the Windows 2000 domain controller both attempt to use port 389 to gain access to Lightweight Directory Access Protocol (LDAP).

Solution

Before you run Exchange Server *5.5* on a Windows 2000 Server domain controller, you must change the Exchange LDAP port number by using the Exchange Server Administrator program under the protocols section. Change the Exchange LDAP port number before you upgrade Windows NT 4.0 to Windows 2000 Server.

Note You need to restart the Exchange *5.5* directory service after changing the port number.

False Alerts, Server Reboots, Services Restart That Don't Exist When Upgrading from Exchange 5.5

You upgrade an Exchange *5.5* server to an Exchange 2000 server, and Exchange 2000 gives false alerts, tries to restart services that don't exist in Exchange 2000, or tries to restart the server.

This occurs when the Exchange *5.5* servers configured to monitor the server you want to upgrade look for services that existed in Exchange *5.5* but that do not exist in Exchange 2000.

Solution

Before you upgrade a server from Exchange 5.5 to Exchange 2000, remove the configuration on other Exchange 5.5 servers that cause them to monitor the server you want to upgrade.

To remove the configuration on Exchange 5.5 servers

1. Put monitoring servers in maintenance mode for the duration of the upgrade to ensure that no monitoring occurs until you bring those servers out of maintenance mode.
2. Upgrade the server.

   Caution When you do this, all monitoring configurations on the newly upgraded server are lost. If that server monitors other servers, it cannot do so after you upgrade it.
3. When the server comes back online, it is configured just as if it is a new Exchange 2000 installation, with the following provisions:
   - No notifications are configured.
   - Status shows that the server is in a critical state if any of the following services are stopped: Web Storage System, Microsoft Exchange message transfer agent (MTA) stacks, Microsoft Exchange Routing Engine, Microsoft Exchange System Attendant, Simple Mail Transport Protocol (SMTP), or World Wide Web Publishing Service.
   - There is no default warning state.
4. Configure notifications, and then configure other Exchange 2000 servers to monitor resources on the upgraded server.

   Note Exchange 5.5 servers should only monitor Exchange 5.5 servers; Exchange 2000 servers should only monitor Exchange 2000 servers.

Setup Fails When Adding a New Server to a Site

You attempt to add a new server to a site and Setup fails.

Solution

To install Exchange correctly on the new server, you must reinstall Exchange or the components that failed.

Caution Information about the new server might be replicated to other servers on the site. To prevent the site's directory from becoming corrupted, remove the new server from the site's directory before you reinstall the server.

To ensure that Exchange is installed correctly on the site

1. On the Administrator program on the remote server that you specified during Setup, delete the server object for the new server.
2. Run Setup again on the new server, and select **Remove All.**
3. Run Setup again on the new server, and reinstall Exchange.

Public Folders Fail to Replicate

Public folder replication does not occur until the Active Directory directory service replicates. If the interval between directory replications is long, Web Storage System might try to replicate public folders before the new site is added to the directory.

Solution

Wait for the directory to replicate the new site information or force directory replication to occur.

To ensure that public folders replicate successfully

1. Click the **General** tab for the server directory.
2. Click **Update Now.**

   Note If you force directory replication to occur, you should do it during a time when the server's performance is not critical and connection costs are low.

URL Not Available During Initial OnCreate Event

You attempt to view a document URL during an initial OnCreate event and the URL is not available. Many workflow applications send notification mail that contains a link (URL) to the current workflow document. However, a new document's permanent URL does not generate until the document is saved for the first time and committed to Web Storage System. The initial OnCreate Workflow event is executed before the document is initially committed; therefore the permanent URL is not available at that time.

Solution

To solve this problem, you should not send notification mail during the initial OnCreate action. Many types of approval applications have an extra step between initial document creation and when the document is subsequently marked as Ready for Approval. The next approver is notified after this extra step.

Problems With Active Directory and Active Directory Connector

Because Exchange 2000 relies on Active Directory and Active Directory Connector (ADC) when communicating with older versions of Exchange, you might need to troubleshoot Active Directory and ADC.

Cannot Modify Active Directory Objects

You attempt to modify Active Directory and cannot complete the task. There are two possible reasons why you cannot modify a directory object:

- The object is read-only.
- You are logged on under an account that does not have permission to modify objects.

Solution

Verify that you have permission to modify objects for the site.

> **Note** Because some objects reference objects in another site, you might also need permission to another site to modify those objects.

Exchange 2000 Options Unavailable in Active Directory Users and Computers

You are using Active Directory Users and Computers and you cannot add, delete, or move an Exchange 2000 mailbox.

This occurs when you administer users where the Exchange System Management tools are not installed. For example, if you install Exchange 2000 Server on a member server and then use Active Directory Users and Computers in Microsoft Management Console (MMC) on a domain controller, you do not see the Exchange 2000 extensions.

Solution One

Start Active Directory Users and Computers from the Exchange 2000 member server.

Solution Two

Install the Exchange 2000 System Management tools on the domain controller.

Hidden Objects in Exchange 5.5 are Visible in Exchange 2000

Exchange Server 5.5 and Exchange 2000 Server are installed in the same site and have an Active Directory connection agreement configured, and you can view replicated hidden objects in Exchange Server 5.5 in the Exchange 2000 System Manager.

Solution

By using the Exchange Server 5.5 Administrator program, you can mark objects as hidden. To view these objects through the Administrator program, on the **View** menu, click **Hidden Recipients.** Exchange 2000 allows you to hide objects from the client view also, but there is no special view in System Manager. Objects that replicate from Exchange Server 5.5 that are marked as hidden remain hidden from the client view but are visible in the Exchange 2000 System Manager.

ADC Creates a Configuration Connection Agreement

You upgrade Exchange Server 5.5 to Exchange 2000 Server, and the installation process requires ADC between Windows 2000 Active Directory and the Exchange Server 5.5 directory service. This creates a type of connection agreement called a configuration connection agreement and homes it on ADC. This connection agreement then reads the site information from Exchange Server 5.5 and writes it into Active Directory configuration naming context.

In addition to having ADC in place, you must also have connection agreements for the Exchange Server 5.5 containers that hold the mailboxes, and any other items you want to upgrade to their Exchange 2000 equivalents. Each connection agreement points to a particular organizational unit in Active Directory.

To access both directories, the connection agreement uses LDAP on port 389. If Exchange 5.5 Server runs on a global catalog or Windows 2000 Server domain controller, you must change the LDAP port of the Exchange 5.5 Directory Service to a value different from the default value.

Solution

Part 1

Check to see if the server is a domain controller.

1. Open Active Directory Users and Computers.
2. Expand the Domain Controller folder.
3. If the server is listed, it is a domain controller. Proceed to Part 2 of this solution.

Part 2

Change the LDAP port in Exchange 5.5.

1. Open the Microsoft Exchange Administrator program.
2. Expand the **Site, Configuration,** and **Protocols** containers, and then double-click **LDAP (Directory)Site Defaults.**
3. On the **General** tab, in the **Port number** box, change the port number value to something other than 389.
4. Click OK.
5. Close the Administrator program and open the MMC Services tool.
6. Restart the directory service.

Active Directory Connector Setup Doesn't Update Schema

You attempt to run Active Directory Connector Setup when you are logged on as a user in the Administrator group, and the operation fails to update the schema.

Solution

To run Active Directory Connector Setup, you must be an Enterprise Administrator. Some operations (during install or reinstall) fail when the user running Active Directory Connector Setup is not an Enterprise Administrator.

Web Storage System Problems

You might have problems with your data storage in Web Storage System, including those with the Extensible File System, mailboxes, and databases.

ExIFS Doesn't Start After Stopping Manually

You stop the Exchange Installable File System (ExIFS) service and then attempt to restart it, and the following error message occurs:

```
System error 2 has occurred.
The system cannot find the file specified.
```

The ExIFS installs with Exchange 2000 Server. The ExIFS runs as a hidden service on an Exchange 2000 server and allows file-level access to various items within a private or public store. On the Exchange 2000 server, the ExIFS service typically maps as drive M in Microsoft Windows Explorer.

Because the ExIFS runs as a hidden service, it is not visible in the **Services** dialog box in **Administrative Tools.** However, you can still issue a **net stop** or a **net start** command to start or stop the ExIFS service. Stopping the ExIFS service does not cause the mapping to drive M to disappear.

Solution

If you manually stop the ExIFS service by issuing a **net stop** command, you must issue the following command at the command line before restarting the ExIFS:

```
Subst m: /d
```

Then you can start the ExIFS service by issuing the **net start** command as shown below:

```
Net start exifs
```

If ExIFS still does not start, verify that no users have a file open on drive M or are accessing drive M by using a share. Alternatively, you can stop all connections to the server by issuing the following command at the command prompt:

```
Netses/ed
```

Caution Issuing the **netses/ed** command will disconnect all users from the server.

Remote File Operation Against an IFS Share Fails With Error 0xEFAD2521

You perform operations on a network drive mapped to an ExIFS share, and the following error occurs after an Exchange 2000 server fails:

```
Error 0xEFAD2521
```

The error might continue to occur even after the server restarts. This error occurs because Exchange does not recognize the mapped network drive.

Solution

To solve this problem, delete the mapped network drive and then re-create it. You can do this at the command prompt by typing the following:

```
net use <drive> /d
net use <drive> \\<server>\<share>
```

"System Cannot Find the Path Specified" or "x:\mailbox is Not accessible" Error

These errors occur in Microsoft Internet Explorer when you attempt to remotely access a mailbox through the MBX share, even though you can successfully access the contents of the MBX directory:

```
System cannot find the path specified
```

```
x:\mailbox is not accessible
```

This happens because the administrator has shared out the MBX directory using a **net share** command and incorrectly typed the domain name.

Solution

Delete the MBX share and then re-create it with the correct domain name.

Last Access Time Is Not Updated On Files

You need to find out the last time a file was accessed, but the **last access time** field has not been updated.

The **last access time** field on files accessed for read-only purposes is disabled by default for performance reasons.

Solution

Enable **last access time** update by creating the following DWORD registry key and set it to 1:

```
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\MSExchangeIS\ParametersSystem\UpdateLastAccessTime
```

Cannot Mount an Additional Web Storage System

You create an additional Web Storage System on your Exchange 2000 server, and you must mount the Web Storage System before it is available for client access. However, when you attempt to mount the store, the following error message occurs:

```
The store could not be mounted because the Active Directory information was not replicated yet.
```

Solution One

Press **Cancel** and mount the store later from its **Context** menu.

Solution Two

Press **Retry** to keep trying to mount Web Storage System. If you click **Retry,** Exchange 2000 tries to mount Web Storage System after the next Active Directory replication cycle, and the following progress bar displays:

```
Waiting for the replication of Active Directory information on server
[your server].
```

Solution Three

To expedite this process, you can manually force Active Directory to replicate.

To manually force Active Directory to replicate

1. In Administrative Tools, open **Active Directory Sites and Services.**
2. Click to expand the following objects:
   - Active Directory Sites and Services ***sewername***
   - Sites
   - Default-First-Site-Name
   - Servers
   - Any domain controller in your Exchange 2000 Server domain
   - NTDS Settings
3. Right–click the object in the details pane, and then click **Replicate Now.**

Exchange 5 to Exchange 2000 Mailbox Move il

You try to move a mailbox from Exchange 5.5 to Exchange 2000, and the following error message occurs:

```
Exchange 5.5 to Exchange 2000 Mailbox Move Failed
```

This occurs because Web Storage System cannot find the specified object. Either the object doesn't have Read permissions or it doesn't fully replicate.

Solution

Either grant the object permission to read Web Storage System or ensure that it fully replicates.

Client Displays Incorrect Public Folders After Modifying Public Database

You modify the Public Store attribute (**msExchHomePublicMDB**) of an Exchange 2000 mailbox to point to Exchange Server *5.5* Public Database, and then modify the Public Store attribute to point back to the Exchange 2000 server, and the Exchange 2000 mailbox still reflects the Exchange Server *5.5* computer.

This occurs because the store provider on the client computer caches the **PR_PROFILE_SERVER** and **PR_PROFILE_SERVER_DN** attributes in the registry as MAPI profile–stored data. This information is updated only for Exchange 2000 to Exchange Server *5.5* but not for Exchange Server *5.5* to Exchange 2000.

Solution

Delete the MAPI profile on the Exchange 2000 client, and then re-create the MAPI profile.

1. Create an Exchange 2000 mailbox.
2. On the Private Store object of the Exchange 2000 mailbox, change the default associated MAPI top-level hierarchy from Exchange 2000 Public Store to the Public Store attribute of an Exchange *5.5* server. The Exchange 2000 mailbox displays the Exchange *5.5* public folders.
3. On the Exchange 2000 mailbox, change from the Exchange Server *5.5* Public Store attribute to the Exchange 2000 Public Store attribute.

The Exchange 2000 mailbox displays the Exchange 2000 public folders after the public mailbox store attribute changes back to reflect the Exchange 2000 server.

Mount or Dismount All Databases Option Missing in a Storage Group

You attempt to mount or dismount all databases in a storage group but there is no option in Exchange 2000 Server System Manager to mount or dismount all databases in a storage group.

Solution

In Exchange 2000, you can have several databases in a storage group. You must mount or dismount each database individually.

To individually mount or dismount each database

1. Start the Exchange 2000 Server System Manager.
2. Click **Organization,** click **Servers,** and then click the name of the server you want.
3. Click **Information Store,** click **storage group,** and then click the database you want to mount or dismount.
4. Right–click the database, click **All Tasks,** and then mount or dismount Web Storage System.

Connectivity Problems

Connectivity problems can include client connectivity problems, protocol problems, problems with public folders, and e-mail and Instant Messaging problems.

Outlook Client Messages Are Lost

You send messages using Exchange 2000 with a Microsoft Outlook 2000 client and message failures occur.

Solution

When you use an Outlook 2000 client with Exchange 2000, you can track messages by using Message Tracking Center, located on the **Tools** menu in System Manager, and Queue Viewer, found on the SMTP and X.400 protocol nodes for each server. To determine the source of a message routing problem between Outlook clients in an Exchange 2000 messaging system, use the flowchart in Figure 33.1 and Table 33.1 to troubleshoot the problem. The event ID codes for each step appear in the appropriate log in Message Tracking Center.

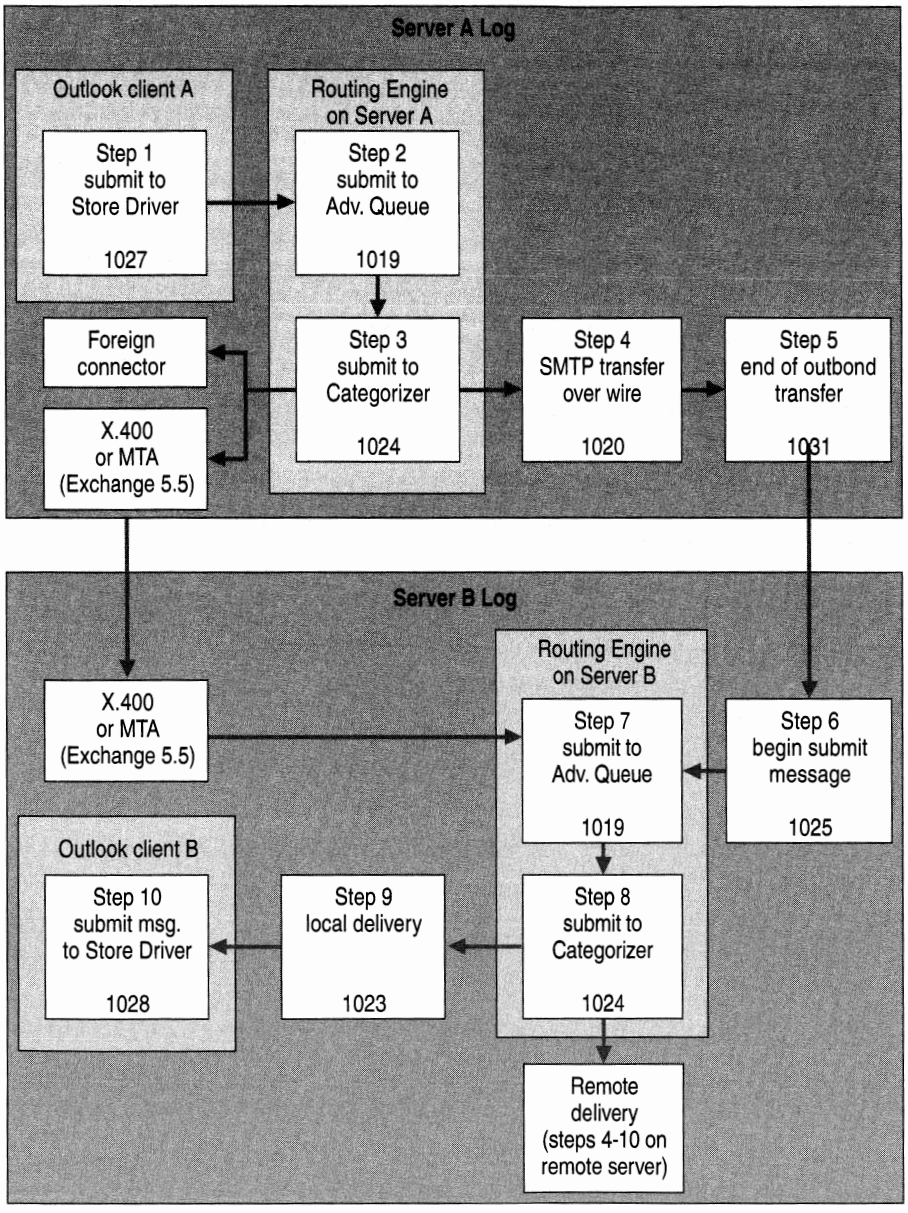


Figure 33.1 Steps an Outlook message takes from server A to sewer B

Note The steps in the server A Log appear on the server that sends the message, whereas the steps in the server B Log appear on the recipient server.

Use Table 33.1 with Figure 33.1 to help troubleshoot Outlook 2000 messaging problems.

Table 33.1 Steps an Outlook message takes

Step	Event ID	Action	Troubleshooting Steps
1	1027	Submit to store driver: message goes from the Outlook 2000 outbox to the store driver	Check to see if the message left the Outlook 2000 outbox.
2	1019	Submit to Advanced Queue: message goes from the store driver to the Advanced Queue	SMTP service might be stopped. Ensure that the SMTP service is running on the sender's home server.
3	1024	Submit to Categorizer: message transfers from the Advanced Queue to the Categorizer	Use the Queue Viewer to see if the message stops in the Advanced Queue. This is where Exchange decides what kind of protocol connector to use. • If SMTP, proceed to Step 4. • If a foreign connector, see documentation for that connector. • If X.400 or MTA (Exchange 5.5), see Exchange 5.5 Help. The message flow rejoins the Exchange 2000 sequence at step 7.
4	1020	SMTP transfer over the wire: message transfers over the wire using SMTP	SMTP service might be stopped. Active Directory might not be functioning. The domain controller might not be functioning. Use Queue Viewer to see if there is a Lookup Queue backlog.
5	1031	End Outbound Transfer	Remote end server might be down.
6	1025	Begin submit message: the message begins transfer over port 25 to the remote server	SMTP service might be stopped. Check to see if it is running.
7	1019	Submit to Advanced Queue: message goes from the store driver to the Advanced Queue	SMTP service might be stopped. Ensure that the SMTP service is running on the remote server.

Table 33.1 Steps an Outlook message takes *(continued)*

Step	Event ID	Action	Troubleshooting Steps
8	1024	Submit to Categorizer: Message transfers from the Advanced Queue to the Categorizer	Use Queue Viewer to see if the message stops in the Advanced Queue.
9	1023	Local Delivery	Web Storage System might be stopped. If the Web Storage System stops, users also have trouble logging on.
10	1028	Submit message to store driver: message goes from SMTP IIS to the store	N/A

Ics.dat File is Corrupted or Lost

Internet News Service creates a file called Ics.dat that is located in the Insdata directory. If the Ics.dat file is corrupted or lost, Internet News Service creates a new Ics.dat file. During the next Network News Transfer Protocol (NNTP) connection, Internet News Service attempts to replicate all the messages in the newsgroup public folders included in the newsfeed.

Solution

To prevent this from occurring, select **Mark All as Delivered** on the **Advanced** tab for that newsfeed.

Site Replication Service Does Not Run in Native Exchange 2000 Environment

You attempt to create a new Site Replication Service when your Exchange 2000 Server organization consists entirely of Exchange 2000 servers, and the following error message occurs:

```
Microsoft Exchange Administrator

This operation cannot be run in a Native Mode Exchange organization. ID
no: c1037d38.
```

Site Replication Service provides support for previous versions of Exchange Server and is not applicable in a native Exchange 2000 environment.

Solution

Do not run Site Replication Service in native mode; run it only in mixed mode.

Instant Messaging Logon Fails in Exchange 2000 but Works in Windows 2000

A user enabled for lnstant Messaging attempts to log on to an Instant Messaging home server and receives the authentication failure dialog box, but the credentials allow the user to log on to Windows 2000. The user is denied access by Exchange 2000 Internet Information Services (IIS).

This occurs because of one of the following conditions:

- The password isn't correct
- The user account is not functioning
- The password policy for the user is not set correctly (enable password encryption)

Solution

To solve this problem, make sure that the password is valid. Make sure that the user account is valid and functioning and the password policy for the user is set correctly.

Instant Messaging Contacts Do Not Display in an Online Conference

You attempt to view Instant Messaging contacts during an online conference and the action fails. In Exchange 2000, Instant Messaging and Conferencing technologies are not integrated; a client's Instant Messaging contacts cannot display in an online conference session.

This occurs because the Microsoft Network (MSN) Messenger client has a user object syntax wherein **User.LogonName** is different than **User.FriendlyName**, which causes an Instant Messaging display problem.

Solution

You can display current Instant Messaging contacts whenever a user joins an online conference by making the following changes in Active Directory:

```
Directory Location
Exchsrvr\Conferencing\Usa

File Name
Confpanel.asp
```

Search the Confpanel.asp file for **User.FriendlyName**, replace it with **User.LogonName**, and save it. After you replace this text, all Instant Messaging client-defined contacts appear in your online conference Web page in the console tree near the bottom.

Instant Messaging Client Disconnects When Changing Status

An Instant Messaging client changes status quickly between one or more different states, and the client automatically disconnects.

This occurs when the Instant Messaging server is busy and rejects client requests such as a request to change status. If a client cannot maintain its status, it autoinatically logs off from the Instant Messaging server.

Solution

This client behavior is by design.

"Conference Not Found" Error Message Displays Using Word 2000 as E-mail Editor

You use Microsoft Word 2000 as your e-mail editor and a "Conference not found" error displays when you attempt to join an online conference. In Exchange 2000, you can join an online conference in a variety of ways. However, when you have Word 2000 defined as your default editor for e-mail messages and you open an e-mail message and click the conference URL, the following error message displays on the Web page:

```
Conference Not Found
The Conference Management Service cannot find the conference you are
trying to join.
```

This occurs because when you use Word 2000 as your default e-mail editor, the URL does not include the equal sign (=), which is added to the end of each conferencing URL.

Solution

Join the conference by clicking the conferencing URL from the **Preview** pane. However, you should not use Word 2000 as your e-mail editor if you join online conferences frequently.

Deleted Newsgroup Remains in Exchange 2000 Public Store

You install Exchange 2000 Server on a computer running Windows 2000 Server, and you have the option to store your newsgroup content in one of three places:

- Local hard disk
- Remote share
- Exchange public folder database

When you store newsgroups in the Exchange public store, you might notice that the corresponding public folder does not delete when you delete the newsgroup in the Exchange System Manager in MMC.

This occurs because Exchange System Manager deletes the newsgroup from the newsgroups list, but the NNTP service does not delete the public folder from the store.

Solution

To delete the public folder, you must to delete it manually from the Public Folders node in Exchange System Manager.

In Exchange System Manager, you can delete newsgroups that exist in Web Storage System. After you expand the NNTP virtual server's node, you can select the Newsgroups node. A list of newsgroups displays in the details pane. You can then select a newsgroup and delete it by clicking **Delete** on the **Action** menu. When you delete a newsgroup in this way, its corresponding public folder is not deleted.

"Deleted Server" Appears in Replication Monitor

When a recently demoted domain controller is replicating, ReplMon might display the following error message:

```
Deleted Server
```

This occurs because Exchange continues to attempt inbound replication from deleted servers for a configurable period of time (the default is two weeks) to allow you to restore domain controllers that are accidentally deleted. After the time period passes, the tombstone is deleted and the deleted server message no longer appears.

Solution

No action is necessary.

Configuring SMTP Connector to Run on tandard Port to Enhance Security Does Not Work

You attempt to configure the SMTP connector to run on a remote port other than port 25 as a means of security to prevent sniffing, but it does not work.

This occurs because Exchange System Manager does not support connection to a remote port other than port 25.

Solution

You should use SMTP authentication or encryption between the two servers configured using Exchange System Manager, and use only port 25 for the SMTP connector.

SMTP Error: "Recipient Could Not Be Reached"

You run Domain Name System (DNS) resolution and the SMTP service causes the following error:

```
The following recipient could not be reached:
foo@server.invalid.com on 5/18/2000 5:26 PM.
The e-mail system was unable to deliver the message, but did not report a
specific reason. Check the address and try again. If it still fails,
contact your system administrator.
<user1-server.user1.extest.microsoft.com #5.0.0>
```

Windows 2000 DNS Server includes the Internet Protocol (IP) addresses of the InterNIC root nameservers preinstalled. This means that a request for a domain that is not defined in a zone on the DNS server is forwarded to one of those servers. If your server is behind a firewall and cannot reach these servers, the "Authoritative Host Not Found" error does not occur. Instead, the "server failed" error occurs.

The SMTP service takes a name such as Remote, which might be an internal fully qualified domain name (FQDN) of a server or an external FQDN of an e-mail domain, and resolves it.

To resolve the FQDN, the SMTP service takes the following steps

1. Exchange checks DNS for a mail exchanger (MX) record for Remote.
2. If DNS returns any entries, it connects to port 25 on each one, in order of lowest priority.
3. DNS returns "Authoritative Host Not Found [1]" if the name server can access the root (.) node of DNS and does not find a record for the domain name.
4. If DNS returns any other error, or no MX entries, it returns to step 2.
5. Exchange calls gethostbyname() for Remote. This results in both an A record search and WINS lookup.

Solution

Make sure the server is configured to use a valid DNS server or servers. Verify that the domain is valid. You can check to see if a domain is valid by using nslookup.

Messages Not Flowing Between Servers in Different Routing Groups

After you create multiple routing groups and move servers into these routing groups, messages no longer flow between servers in the different routing groups. These messages are immediately returned to their sender with a 5.0.0 error code in the non-delivery report (NDR).

This occurs because there must be connectors configured between routing groups for mail to flow between them.

Solution

Configure Routing Group connectors for the routing groups.

SMTP Protocol Error "454 Client Must Be Authenticated"

Clients are having trouble sending e-mail to an Exchange 2000 server and the following SMTP protocol error occurs:

```
454 Client must be authenticated
```

This is because the server does not allow anonymous SMTP connections and requires the client to authenticate. SMTP clients that have not been configured for SMTP authentication or do not have this capability cannot submit e-mail to a server that is set to disallow anonymous connections.

Solution One

Configure the SMTP client to perform SMTP authentication to the server.

Solution Two

Configure the Exchange 2000 server to allow anonymous connections.

Cannot Gain Access To Additional HTTP Virtual Server on Cluster

You create an additional HTTP virtual server (using Exchange System Manager) in a Windows 2000 Cluster server, but the newly created HTTP virtual server is not accessible.

This occurs because Exchange 2000 Setup does not create the correct **HostName** for Distributed Authoring and Versioning (DAV).

Solution

Do not create additional HTTP virtual servers in an active/passive Microsoft Exchange 2000 Server cluster environment.

Windows 2000 Server Tools Problems

Exchange 2000 relies on a number of Windows 2000 Server features and tools to operate. You might find it necessary to troubleshoot some of the Windows 2000 Server tools.

Cannot Access Registry

You encounter an error message that indicates Exchange cannot access the registry because an incorrect value is detected in this single instance and the registry cannot update.

Solution

Delete the instance and then add it back. Deleting the instance removes the old instance and allows you to start over.

Deleting an instance that has a corrupt registry entry might not be successful. You can continue by performing the following procedure.

> **Caution** Do not use a registry editor to edit the registry directly unless you have no alternative. The registry editor bypasses the standard safeguards provided by administrative tools. These safeguards prevent you from entering conflicting settings or settings that are likely to degrade performance or damage your system. Editing the registry directly can have serious, unexpected consequences that can prevent the system from starting and require that you reinstall Exchange 2000. To configure or customize Exchange 2000, use the programs in Control Panel or MMC whenever possible.

To delete a registry instance

1. Stop the external instance in the Control Panel\Services dialog box.
2. Remove the instance name from the following registry value:

   ```
   HKLM\Software\CurrentControlSet\NT External\Linkage\Export" MULTI-SZ
   ```

3. Delete the following instance key from the registry:

   ```
   HKLM\Software\CurrentControlSet\<Instance Name>
   ```

This removes the external instance completely.

Backup and Restore Problems

Backup and restore are important functions for maintaining Exchange 2000. You might find it necessary to troubleshoot these functions.

Cannot Create Identical Mailboxes in a Site or Organization

You attempt to create a copy of an Exchange 2000 server within the same organization as the original for backup or restore purposes, and Exchange 2000 does not function properly.

This occurs because you cannot create an identical server or identical mailboxes in a site or organization for backup. Exchange 2000 requires a unique Exchange distinguished name.

Solution

This is by design

Event ID 1018,1019, and 1022: Database is Damaged

When performing backup and restore, you receive an Event ID 1018, 1019, or 1022: Database is damaged error message. This means that online backup cannot complete because the database is damaged.

Solution

Check the hardware for errors and complete a restore of this database as soon as possible. You can perform an offline backup so you have a recent copy of the database on tape even though it is damaged.

> **Caution** You should never delete logs from the system when performing an offline backup. The logs are required if you want to restore from the online backups.

Storage Group Fails to Mount with -1216 (JET_errAttachedDatabaseMismatch) or 0xfffffb40 Error

When mounting a storage group with a missing database in an attempt to recover after an unexpected shutdown, you receive a -1216 error (from ESE98) or an 0xfffffb40 error in the application log.

This is because the Extensible Storage Engine (ESE) attempts to bring all databases in a storage group to a consistent state during recovery. To accomplish this, ESE keeps track of all databases in the log files for the storage group. If a database is missing, ESE will return the error -1216 and will not start the storage group.

Solution One

Find the missing database files and place them in the correct locations.

Solution Two

If the missing database was deleted or lost, restore the missing database from backup. While restoring this database, the other databases in the storage group can be mounted and accessed by completing the following steps:

1. Run **eseutil /r** ***Log Base Name*** **/l** to allow ESE to mount the databases that are present while ignoring the missing databases. After running this command, any databases that are not present will need to be restored from backup.
2. Mount the databases in the storage group that are now present and consistent.
3. Allow the users on the running databases to access their e-mail.
4. Restore the lost database from backup.

Performance Problems

If you identify a resource that functions outside the recommended performance threshold, you should investigate that resource in greater detail. This might include the following steps:

- Analyze your hardware and software configurations. Your configuration should match Microsoft recommendations for the operating system and the services you support.
- Review entries in the event log for the time period when you begin seeing out-of-range counter values; these entries might provide information about problems that can result in poor system performance.
- Examine the applications you are running and what resources they require, to determine their adequacy.
- Consider variables in your workload, such as the possibility of processing different jobs at different times. For more efficient analysis, when looking for a specific problem, limit your charts and reports to specific events occurring at known times.
- For immediate diagnosis and problem solving of situations such as shutdowns and logon failures, log or monitor for a shorter time period. You should sample frequently when monitoring over a short period. Similarly, for long-term planning and analysis, you should log for a longer period and set the update interval accordingly.

- Consider network or hard disk use or other activities that occur at the times when you see increased resource utilization. Try to understand the usage patterns. They might be associated with specific protocols or computers.
- Approach corrections in a scientific manner. For example, never make more than one change at a time, always repeat monitoring after a change to validate the results, eliminate results that are suspect, and keep records of what you do and what you learn.

When investigating slow downs (also known as "bottlenecks" in specific resources), focus on the performance objects and counters that pertain to the specific resource where you think the problem resides.

Windows 2000 System Monitor provides information about performance statistics such as processor use, memory use, server throughput, queue size, and number of write/read operations per second. You can use the pre-configured System Monitor chart views that are included with Exchange 2000 to help you improve your system's performance and identify bottlenecks.

Memory Leak in LDAP Service When Using Migration Wizard

You use LDAP Migration Wizard and perform multiple scans on large containers, and the LDAP service does not release memory. For each container scanned, the service adds more memory. Eventually all available memory is consumed and CPU usage is at 100 percent.

Solution

When memory gets low, a message from Windows 2000 Server warns of low virtual memory and asks you to close some applications. Closing applications alleviates the memory problem.

APPENDIX B

Ports and Protocols

You should review the ports and protocols listed in this appendix to prevent interruption of service when implementing a firewall.

The following is a list of ports and protocols for Microsoft Windows 2000 services.

Port	TCP/UDP	Service Name
42	TCP	WINS Replication
47	TCP	GRE for PPTP
53	UDP	DNS Name Resolution
53	TCP	DNS
67	UDP	DHCP Lease (BOOTP)
68	UDP	DHCP Lease
88	UDP	Kerberos
135	TCP	Location Service (RPC, RPC EP Mapper, WINS Manager, DHCP Manager, MS DTC)
137	UDP	NetBIOS Name Service (Logon Sequence, Windows NT 4.0 Trusts, Windows NT 4.0 Secure Channel, Pass Through Validation, Browsing, Printing)
137	TCP	WINS Registration
138	UDP	NetBIOS Datagram Service (Logon Sequence, Windows NT 4.0 Trusts, Windows NT 4.0 Directory Replication, Windows NT 4.0 Secure Channel, Pass Through Validation, NetLogon, Browsing, Printing)
139	TCP	NetBIOS Session Service (NBT, SMB, File Sharing, Printing, Logon Sequence, Windows NT 4.0 Trusts, Windows NT 4.0 Directory Replication, Windows NT 4.0 Secure Channel, Pass Through Validation, Windows NT 4.0 Administration Tools [Server Manager, User Manager, Event Viewer, Registry Editor, Diagnostics, Performance Monitor, DNS Administrator])
389	TCP/UDP	LDAP

(continued)

Port	TCP/UDP	Service Name
500	TCP/UDP	ISAKMP/Oakley negotiation traffic (IPSec)
522	TCP	User Location Store
636	TCP/UDP	LDAP (over TLS/SSL)
750	UDP	Kerberos Authentication
750	TCP	Kerberos Authentication
751	UDP	Kerberos Authentication
751	TCP	Kerberos Authentication
752	UDP	Kerberos Password Server
753	UDP	Kerberos User Registration Server
754	TCP	Kerberos Slave Propagation
888	TCP	Logon and Environment Passing
Dynamic	TCP	Directory Replication
1109	TCP	POP with Kerberos
1723	TCP	PPTP Control Channel (IP Protocol 47 – GRE)
2053	TCP	Kerberos de-multiplexor
2105	TCP	Kerberos encrypted login
3268		Global Catalog
3269		Global Catalog
3389	RDP	Terminal Services

The following is a list of ports and protocols for Microsoft Exchange 2000 Server services.

Port	TCP/UDP	Service Name
25	TCP	SMTP
80	TCP	HTTP
102	TCP	MTA – X.400 over TCP/IP
110	TCP	POP3
119	TCP	NNTP
135	TCP	Client/Server Communication, RPC, Exchange Administration
143	TCP	IMAP4
389	TCP	LDAP
443	TCP	HTTP (SSL)
465	TCP	SMTP (SSL)
563	TCP	NNTP (SSL)
636	TCP	LDAP (SSL)
993	TCP	IMAP4 (SSL)
995	TCP	POP3 (SSL)
1720	TCP	H.323 Call Setup
1731	TCP	Audio Call Control
2980	TCP/UDP	Instant Messaging Service
Dynamic	TCP	H.323 Call Control
Dynamic	UDP	H.323 Call (RTP Over UDP)

Glossary

access control
The security mechanism in Windows 2000 that limits access to information, objects, or controls for designated users and groups.

access control entry
(ACE) An entry in an access control list (ACL) that contains the security identifier (SID) for a user or group and an access mask that allows, denies, or audits operations by users or groups.

access control list
(ACL) A list of Windows 2000 security principals, user accounts, and groups associated with an object. This list is used to determine whether a user or process has been granted access to an object.

access token
A data structure containing security information that identifies a user to the security subsystem on a computer running Windows 2000 or Microsoft Windows NT. Access tokens contain a user s security ID, the security IDs for groups that the user belongs to, and a list of the user's privileges on the local computer.

ACE
See definition for: access control entry

ACL
See definition for: access control list

Active Directory
The directory service for Windows 2000 Server. It stores information about objects on the network and makes this information available for authorized administrators and users. Active Directory gives network users access to permitted resources anywhere on the network using a single logon process. It provides administrators with an intuitive hierarchical view of the network and a single point of administration for all network objects.

Active Directory Connector
(ADC) A Windows 2000 service that replicates the Exchange *5.5* directory with Active Directory. This allows administration of a directory from either Active Directory or the Exchange *5.5* directory service.

Active Directory Service Interface
(ADSI) A set of interfaces that allows programmatic access to underlying directory services through a common command set.

Active Directory Users and Computers
A Microsoft Management Console (MMC) snap-in that allows administrators to manage objects in Active Directory.

Active Server Pages
(ASP) A scripting environment that runs ActiveX scripts and ActiveX components on a server. Developers can combine scripts and components to create Web-based applications.

ADC

See definition for: Active Directory Connector

address list

A collection of recipient and other Active Directory objects. Each address list can contain one or more types of objects (for example, users, contacts, groups, public folders, conferencing, and other resources). Exchange 2000 address lists also provide a mechanism to partition mail-enabled objects in Active Directory for the benefit of specific groups of users.

address space

A set of address information associated with a connector or gateway that identifies certain types of messages. An address space is typically a subset of a complete address.

administrative group

A collection of Active Directory objects that are grouped together for the purpose of permissions management. An administrative group can contain policies, routing groups, public folder hierarchies, servers, and chat networks. The content of an administrative group depends on choices you make during installation.

ADSI

See definition for: Active Directory Service Interface

advanced security

A feature that enables users to digitally sign or encrypt messages. To sign a message, the sender must provide an advanced security password. This guarantees recipients that a digitally signed message is authentic. To decrypt a message, recipients must provide an advanced security password.

ASP

See definition for: Active Server Pages

asynchronous event

An event that occurs after an item is saved or deleted. An asynchronous event does not occur in any particular sequence with other events.

attribute

1. Information that indicates that a file is read-only, hidden, system, or compressed, or whether the file has been changed since a backup copy of it was made. 2. In object-oriented software, an individual characteristic of the object.

audit

To track the activities of users by recording selected events in an event log on a server or workstation.

authentication

1. Validation of a user's Windows 2000 logon information. 2. The process that verifies the identity of a user trying to establish a connection to a chat server. Chat Service supports authentication of incoming client connections by using cleartext passwords, NTLM protocol, or any authentication method compatible with the Security Support Provider Interface (SSPI).

B

back-end server

A server that hosts at least one database that front-end servers connect to when relaying requests from clients.

See also: front-end server

backbone

The network connection between LAN segments.

ban
A control that allows users and administrators to restrict users with a specific user name or nickname, or users from a specific domain from participating in a chat community.

bastion host
A computer that must be secure because it is accessible from the Internet and exposed to attack. It acts as a protective relay for mail between the Internet and internal users.

bridgehead server
A computer that connects servers using the same communications protocols so information can be passed from one server to another. In Exchange 2000, a bridgehead server is a connection point from a routing group to another routing group, remote system, or other external system.

C

CA
See definition for: certification authority

CDO
See definition for: Collaboration Data Objects

certificate
An electronic credential that authenticates a user on the Internet and intranets. Certificates ensure the legitimate online transfer of confidential information or other sensitive material by means of public encryption technology. In Exchange, certificates contain information used for digital signatures and encryption that binds the user s public key to the mailbox.

certificate revocation list
(CRL) The list of users who have had their security tokens revoked, and therefore should not be authenticated as secure.

Certificate Services
Software services that provide authentication support including secure e-mail, Web-based authentication, and smart card authentication. The services contrast with Internet Authentication Services (IAS), which provide authentication for dial-in users.

certificate template
A Windows 2000 construct that pre-specifies the format and content of certificates based on their intended usage.
See also: public key infrastructure

certificate trust list
(CTL) A signed list of root certification authority certificates that an administrator considers reputable for designated purposes, such as client authentication or secure e-mail.

certification authority
(CA) An entity with a server that issues certificates to clients and servers. A certification authority attests to the identification of a user of a public key and can also revoke certificates when the private key associated with the certificate is compromised or when the subject of the certificate leaves an organization.

channel
A channel, also called a chat room, is the Chat Service platform for communication. When users join a channel they can read anything that is typed to the members of the channel.

checkpoint file
A file that indicates which transactions have been successfully saved to disk. The Edb.chk file points to the log file of all transactions that have been successfully committed to the database file. After all the transactions in a particular log file are committed to the database file, the pointer advances to the log file with the next unwritten entry. Separate checkpoint files are maintained for each storage group.

circular logging
A method of logging transactions in Microsoft Web Storage System in which earlier log files are overwritten after the transactions in the log file have been committed to the database.

coexistence
When you connect Exchange 2000 to another messaging system, including an earlier version of Exchange, the two systems coexist. A coexistence period can be short-term (enough time to migrate users from an existing messaging system to Exchange 2000), or it can be long-term (a permanent connection to the messaging system of another department that is not moving to Exchange 2000).

Collaboration Data Objects
(CDO) An application programming interface that allows users and applications high-level access to data objects within Exchange. CDO defines the concept of different object classes, including messages, posts, appointments, and tasks.

Conference Management Service
The component of Exchange 2000 Conferencing Server responsible for the reservation and scheduling of online meetings.

See also: Exchange 2000 Conferencing Server

conference resource
An Exchange 2000 mailbox that users invite when scheduling an online meeting.

configuration connection agreement
A connection agreement that replicates Exchange-specific configuration information between Exchange 5.5 and Active Directory. This type of connection agreement is created automatically the first time an Exchange 2000 server is introduced into an Exchange 5.5 site. It cannot be created manually.

See also: connection agreement

connection agreement
Used by Active Directory Connector (ADC) to control replication between an Exchange 5.x or earlier site and Active Directory. The standard connection agreement replicates Exchange recipient objects (mailboxes, distribution lists, custom recipients, and public folder proxies) and Active Directory objects (users, groups, contacts, and public folder proxies) between the Exchange 5.5 directory and Active Directory. Connection agreements define the server names to be contacted for replication, the object classes to replicate, the target containers, and the replication schedule.

See also: configuration connection agreement, primary connection agreement

console
A control unit through which a user communicates with a computer via a primary input device (keyboard or mouse) and a primary output device (screen). A console integrates all the tools, information, and Web pages an administrator needs to perform specific tasks.

contact
An Active Directory object that represents a user who does not have a Windows logon account or a mailbox. For example, a contact may represent a user outside of the organization. A contact in Windows 2000 is equivalent to a custom recipient in earlier versions of Exchange.

See also: custom recipient

container
An object that contains other objects

contiguous namespace
A namespace that contains names that share a common root; for example, microsoft.com and exchange.microsoft.com form a contiguous namespace.

CRL
See definition jor: certificate revocation list

CTL
See *definition* ***for:*** certificate trust list

custom address list
An address list created for users who need a custom view of recipients within an Exchange organization. For example, you can create an address list that includes only employees in North America, or you can create an address list that includes only employees in the marketing department.

See also: default address list

custom recipient
Used in previous versions of Exchange, a custom recipient was a user who was not hosted by Exchange. In Exchange 2000, such users can be added to Active Directory as contacts, Windows 2000 users, or users whose Windows 2000 accounts are disabled. They are mail-enabled. but not mailbox-enabled, because their mailboxes are hosted on another messaging system.

See also: contact, mail-enabled

D

data conference
Online conferences in which members share data in real time.

Data Conferencing Provider
A conference technology provider supplied with Exchange 2000 Conferencing Server that permits the hosting of data conferences.

default address list
An address list that is automatically created based on the values of specific attributes of Active Directory objects. These address lists are available to Exchange users without any administrator action.

See also: custom address list

DHCP
See definition for: Dynamic Host Configuration Protocol

digital signature
A personal authentication method based on encryption and secret authorization codes that is used for signing electronic documents. Digital signatures not only validate the sender's identity, they ensure the message contents have not been altered. No one can tamper with a digitally signed message without detection. When the sender encrypts a message, only the recipient is able to decrypt it and read its contents.

directory partition
A self-contained section of a directory hierarchy that can have its own properties, such as replication configuration. Active Directory includes the domain, configuration, and schema directory partitions.

See also: naming context

directory replication
The process of updating the directories of all servers within and between sites.

Distributed Authoring and Versioning
(DAV) An extension to the Hypertext Transfer Protocol (HTTP) 1.1 protocol that allows for manipulation of objects and attributes. Although not specifically designed for the purpose, DAV allows for the control of a filing system by using HTTP protocol.

distribution list
A group of recipients created to expedite mass mailing of messages and other information. When e-mail is sent to a distribution list, all members of that list receive a copy of the message.

See also: group

DNS
See definition for: Domain Name System

domain
A group of computers that are part of a network and share a common directory database. In Windows 2000, a domain is a security boundary and permissions that are granted in one domain are not carried over to other domains.

domain controller
A computer running Windows 2000 Server that manages user access to a network, which includes logging on, authentication, and access to Active Directory and shared resources.

domain controller locator
An algorithm that runs in the context of the Net Logon service and that finds domain controllers on a Windows 2000 network. Locator can find domain controllers by using either DNS or network basic input/output system (NetBIOS) names, or it can be used on a network where Internet Protocol (IP) transport is not available.

domain local group
A Windows 2000 group available only in native-mode domains, which can contain members from anywhere in the forest, in trusted forests, or in a trusted pre Windows 2000 domain. Domain local groups can grant permissions only to resources within the domain in which they exist. Typically, domain local groups are used to gather security principals from across the forest to control access to resources within the domain.

See also: universal group

Domain Name System
(DNS) A TCP/IP standard name service that allows clients and servers to resolve names into Internet Protocol (IP) addresses and vice versa. Dynamic DNS in Windows 2000 enables clients and servers to automatically register themselves without the need for administrators to manually define records.

domain naming master
A domain controller that can add new domains to the forest, remove existing domains from the forest, and add or remove cross-reference objects to external directories. Only the domain naming master can perform those tasks.

dual key pair system
A security architecture that uses two separate key pairs, each with separate usage restrictions. One key pair is used for message encryption, while the other is used for generating and validating digital signatures. Exchange Key Management Service uses a dual key pair design so it can archive, and provide recovery of, the user s private encryption key. To prevent the possibility of signature forgery by an administrator, the private signature key is kept solely in the possession of the user.

Dynamic Host Configuration Protocol
(DHCP) A protocol for assigning Internet Protocol (IP) addresses to computers and other devices on a TCP/IP network. Dynamic addressing permits a computer to have a different address each time it logs on to a network.

E

encryption
An advanced security feature that provides confidentiality by allowing users to conceal data. Data is encrypted while it resides on disk and travels over a network.

event
In the context of programming, the occurrence of some particular action or the occurrence of a change of state that can trigger an event sink. For example, the arrival of a message to the SMTP service is an event that can trigger any number of event sinks.

event sink
A piece of code that activates upon a defined trigger, such as receiving a new message. The code is normally written in any COM-compatible programming language, such as Microsoft Visual Basic, Microsoft Visual Basic Scripting Edition (VBScript), JavaScript, C or C++. Exchange 2000 supports the transport, protocol, and store event sinks. Event sinks on the store can be synchronous (code executes as the event is triggered) or asynchronous (code executes sometime after the event).

Exchange 2000 Conferencing Server
An application that provides scalable, reliable online data and video conferences.

See also: Conference Management Service

Exchange Administrator
An Exchange Administration Delegation Wizard role that grants the users permission to fully administer Exchange system information, but not modify permissions.

Exchange Full Administrator
An Exchange Administration Delegation Wizard role that grants the user permission to fully administer Exchange system information and modify permissions.

Exchange View Only Administrator
An Exchange Administration Delegation Wizard role that grants the user permission to view Exchange configuration information.

extended Instant Messaging address
A more fully qualified Instant Messaging address, based on the standard Simple Mail Transfer Protocol (SMTP) format someone@im.microsoft.com.

See also: lnstant Messaging address

Extensible Storage Engine
(ESE) Formerly known as JET, Extensible Storage Engine is a method that defines a very low-level application programming interface (API) to the underlying database structures in Exchange. Extensible Storage Engine is also used by other databases, such as the Active Directory database. The Extensible Storage Engine uses a balanced-tree (B-tree) structure to store data. Each page in the database file is a node in the B-tree structure. An Extensible Storage Engine database can contain up to 2^32 pages or *16* terabytes (Active Directory uses 8 KB pages and can contain up to *32* terabytes).

failover
The process of taking resources, either individually or in a group, offline on one node and bringing them back online on another node.

firewall
A security system intended to protect an organization's network against external threats coming from another network, such as the Internet. A firewall prevents direct communication between an internal network and external computers by routing communication through a proxy server that exists outside the network.

forest
One or more domain trees that do not form a contiguous namespace. Forests allow organizations to group divisions that operate independently but still need to communicate with one another.

FQDN
See definition for: fully qualified domain name

front-end and back-end architecture
An Exchange architecture in which clients access a set of protocol servers (the front end) for collaboration information, and these servers in turn request data from separate servers (the back end). A front-end and back-end architecture provides a scalable, single point of contact for all data requests.

front-end server
A server that receives requests from clients and relays them to the appropriate back-end server.

See also: back-end server

full-text indexing
An indexing feature that allows users to use Microsoft Outlook Advanced Find and custom clients to quickly locate mail messages and documents in Microsoft Web Storage System. Full-text indexing includes message properties, body text, and attachments.

fully qualified domain name
(FQDN) A DNS domain name that has been stated unambiguously to indicate with certainty its location in the domain namespace tree. Fully qualified domain names differ from relative names in that they typically are stated with a trailing period (.), for example, host.example.microsoft.com, to qualify their position to the root of the namespace.

GAL
See definition for: global address list

global address list
(GAL) A list containing all Exchange users, contacts, groups, conferencing resources, and public folders in an organization. This list is retrieved from the global catalog servers in Active Directory and is used by Outlook clients to address messages or find information about recipients within the organization.

global catalog
A server that holds a complete replica of the configuration and schema naming contexts for the forest, a complete replica of the domain naming context in which the server is installed, and a partial replica of all other domains in the forest. The global catalog is the central repository for information about objects in the forest.

global group
For Windows 2000 Server, a group that can be used in its own domain, in member servers and in workstations of the domain, and in trusting domains. In all those places a global group can be granted rights and permissions and can become a member of local groups. However, a global group can contain user accounts only from its own domain.

group
A collection of users, computers, contacts, public folders, and other groups. Groups can be used as a security identifier or as a distribution list. Distribution groups are used only for e-mail. Security groups are used to grant access to resources. A group in Windows 2000 is roughly equivalent to a distribution list in Exchange 5.5.

See also: distribution list

Group Policy
The Windows 2000 Microsoft Management Console (MMC) snap-in used to specify the behavior of users' desktops.

See also: Group Policy object

Group Policy object
A collection of Group Policy settings. Group Policy objects are essentially the documents created by the Group Policy snap-in, a Windows 2000 utility. Group Policy objects are stored at the domain level, and they affect users and computers contained in sites, domains, and organizational units.

***See** also:* Group Policy

H

hash
A fixed-size result obtained by applying a one-way mathematical function (called a hash or message digest function) to an arbitrary amount of data. Given a change in the input data, the resulting hash changes. A hash is also called a message digest.

HTTP
See definition for: Hypertext Transfer Protocol

Hypertext Transfer Protocol
(HTTP) A client/server protocol used on the Internet for sending and receiving HTML documents. HTTP is based on the TCP/IP protocol.

IFS
See definition for: Installable File System

11s
See definition for: Internet Information Services

IMAP4
See definition for: Internet Message Access Protocol

in-place upgrade
A method of upgrading to Exchange 2000 Server in which you run Exchange 2000 Setup on a server that is running Exchange Server 5.5 with Service Pack 3.

infrastructure master
A domain controller that updates cross-domain group-to-user references to reflect a user s new name. The infrastructure master updates these references locally and uses replication to bring all other replicas of the domain up to date. If the infrastructure master is unavailable, these updates are delayed.

Installable File System
(IFS) A storage technology that functions as a filing system. It makes mailboxes and public folders available as traditional folders and files through standard Microsoft Win32 processes, such as Microsoft Internet Explorer and the command prompt.

Instant Messaging address
In Instant Messaging, the use of standard Simple Mail Transfer Protocol (SMTP) e-mail addresses as Instant Messaging addresses. Some situations require the use of a more fully qualified Instant Messaging address, referred to as the extended Instant Messaging address.

See also: extended Instant Messaging address

Instant Messaging domain
A DNS name that identifies a logical collection of Instant Messaging user accounts and home servers represented by a virtual server called an Instant Messaging router. It is recommended that Instant Messaging domains have a one-to-one correspondence with e-mail domains.

Instant Messaging home server
A virtual server that hosts Instant Messaging user accounts and communicates directly with clients to send and deliver instant messages and presence information.

Instant Messaging router
An Instant Messaging router that receives incoming messages, locates the recipient's home server. and forwards the message to that server for delivery to the recipient.

Instant Messaging Service
A service that allows for real-time messaging and collaboration between users.

Internet Information Services
(IIS) Microsoft's Web service for publishing information on an intranet or the Internet, and for building server-based Web applications. Upon installation, Exchange 2000 extends the messaging capabilities of IIS and incorporates them into the Exchange message routing architecture.

Internet Key Exchange
A protocol that establishes the security association and shared keys necessary for two parties to communicate with Internet Protocol security (IPSec).

Internet locator service
Active Directory uses DNS as an internet locator service, resolving Active Directory domain, site, and service names to an IP address.

Internet Message Access Protocol
(IMAP) An Internet messaging protocol that enables a client to access mail on a server rather than downloading it to the user's computer. IMAP is designed for an environment where users log on to the server from a variety of different workstations.

Internet service provider
(ISP) A business that supplies Internet connectivity services to individuals, businesses, and other organizations.

ISP
See definition for: Internet service provider

K

KDC
See definition for: Key Distribution Center

Kerberos V5
An authentication protocol used to verify user or host identity. Kerberos V5 authentication protocol is the default authentication service for Windows 2000.

key
A code or number used to digitally sign and encrypt data for security-enabled users. Keys often occur in pairs; for example, a public key and a private key. Public keys are issued by third-party certification authorities (CAs). A certificate binds a user s public key to his or her mailbox.

Key Distribution Center
(KDC) A network service that supplies session tickets and temporary session keys used in the Kerberos authentication protocol. In Windows 2000, the KDC runs as a privileged process on all domain controllers. The KDC uses Active Directory to manage sensitive account information such as passwords for user accounts.

Key Management server
The Exchange computer on which the Key Management Service has been installed. There can be one Key Management server per administrative group.

Key Management Service
An optional Microsoft Exchange 2000 Server component that is installed on a designated server in an administrative group. It provides centralized administration and archival of private keys, and maintains every user s private encryption key in an encrypted database. The keys are used for encrypting e-mail messages and signing messages with digital signatures.

key pair
Used in message security, a cryptographic key pair consists of a public key and a private key. A public key is associated with a user through a certificate that is published to a location available to anyone. The corresponding private key is stored in a secure location on the user's client computer. Key Management servers generate key pairs for encryption in Exchange 2000, while Microsoft Outlook generates key pairs for digital signatures.

See also: public key infrastructure

Knowledge Consistency Checker
A built-in process that runs on all domain controllers and generates the replication topology for the Active Directory forest. At specified intervals, the Knowledge Consistency Checker reviews and makes modifications to the replication topology to ensure propagation of data either directly or transitively.

L

LDAP
See definition for: Lightweight Directory Access Protocol

LDAP Data Interchange Format
(LDIF) A draft Internet standard for a file format that can be used to perform batch operations on directories that conform to Lightweight Directory Access Protocol (LDAP) standards.

LDIF
See definition for: LDAP Data Interchange Format

leapfrog upgrade
A method of upgrading to Exchange 2000 Server in which Exchange 2000 is installed on a new server, users are moved from a server running an earlier version of Exchange to the new server, and Exchange is then installed on the server running the earlier version of Exchange. This move and upgrade cycle repeats until all servers are upgraded.

Lightweight Directory Access Protocol
(LDAP) A network protocol designed to work on TCP/IP stacks to extract information from a hierarchical directory such as X.500. It is useful for searching through data to find a particular piece of information.

link state algorithm
The algorithm used to propagate routing status information between Exchange 2000 servers.

link state information
Infonnation about the state of messaging routes (links) in an Exchange 2000 messaging system derived from the link state algorithm to quickly and frequently calculate the state of system links for up-to-date status about routes. Exchange 2000 servers use link state information to make the best routing choice at the source rather than sending a message down a path where a link is not working. This eliminates message bounce and looping.

link state table
The database on each Exchange 2000 server used to store link state information propagated by the link state algorithm. The link state table is used to evaluate cost and availability information to determine the most suitable route for a message.

local bridgehead server
A server within a routing group that handles e-mail flow to and from a connector in that routing group. Routing group connectors can have multiple local bridgehead servers or no local bridgehead server, in which case every server in the routing group acts as a local bridgehead server. SMTP and X.400 connectors must have one, and only one, local bridgehead server.

Local Security Authority
(LSA) A protected subsystem that authenticates and logs users onto the local system. In addition, the LSA maintains information about all aspects of local security on a system (collectively known as the local security policy), and provides various services for translation between names and identifiers.

LSA
See definition *for:* Local Security Authority

M

mail exchanger resource record
(MX resource record) A Domain Name System (DNS) record that specifies a mail exchange server for a DNS domain name. A mail exchange server is a host that either processes or forwards mail for the DNS domain name. Processing the mail means either delivering it to the addressee or passing it to a different type of mail transport. Forwarding the mail means sending it to its final destination server, sending it using Simple Mail Transfer Protocol (SMTP) to another mail exchange server that is closer to the final destination, or queuing it for a specified amount of time.

mail-enabled
An Active Directory object that has at least one e-mail address defined. If the user is mail-enabled, the user has an associated e-mail address, but does not have an associated Exchange mailbox.
See also: custom recipient

mailbox
The location where e-mail is delivered. The administrator sets up a mailbox for each user. If a set of personal folders is designated as the e-mail delivery location, e-mail is routed from the mailbox to this location.

mailbox store
The part of Microsoft Web Storage System that maintains data in mailboxes. A mailbox store consists of a rich-text .edb file, plus a streaming native Internet content .stm file.

mailbox-enabled
An Active Directory object that has an Exchange mailbox associated with it; therefore it can both send and receive messages within the Exchange system.

message transfer agent
(MTA) An Exchange component that routes messages to other Exchange MTAs, information stores, connectors, and third-party gateways. Also referred to as X.400 protocol in Exchange 2000 System Manager.

metabase
A store that contains metadata, such as that used by Internet Information Services (IIS). The metabase can be viewed through utilities such as Metaedit.

Microsoft Management Console
(MMC) A management display framework that hosts administration tools and applications. Using MMC you can create, save, and open collections of tools and applications. Saved collections of tools and applications are called consoles.

Microsoft Security Service Provider Interface
(SSPT) The API for obtaining integrated security services for authentication, message integrity, message privacy, and secure quality of service for any distributed application protocol.

migration
The process of moving an existing messaging system to another system by copying the existing mailboxes, messages, and other data, and importing that information into a new messaging system.

MIME
See *definition* for: Multipurpose Internet Mail Extensions

mixed mode
The default operating mode of Exchange when it is installed. Mixed mode allows Exchange 2000 servers and servers running earlier versions of Exchange to coexist in the same organization. Mixed mode allows interoperability between versions by limiting functionality to features both products share.

MMC
See definition for: Microsoft Management Console

moderated channel
A chat channel that is used for small chat events. A chat user joining a moderated channel cannot post messages to the channel without permissions, but can see messages posted by the designated speakers. A speaker, acting as a channel host, can grant speaking permission to a specific user.

move mailbox upgrade
A method of upgrading to Exchange 2000 Server in which Exchange 2000 is installed on a new server, users are moved from a server running an earlier version of Exchange to the new server, and the server running the earlier version of Exchange is removed.

MTA
See definition ***for:*** Message Transfer Agent

Multipurpose Internet Mail Extensions
(MIME) A standard that enables binary data to be published and read on the Internet. The header of a file with binary data contains the MIME type of the data; this informs client programs (such as Web browsers and mail packages) that they cannot process the data as straight text.

namespace
A set of names associated with a domain or forest that identifies objects that belong to the domain or forest. A DNS name creates a namespace: for example, microsoft.com.

naming context
A term used in X.500 and LDAP standards.

See also: directory partition

native mode
An operating mode of Exchange 2000 Server when it is running only Exchange 2000 Server. Servers running earlier versions of Exchange cannot join an organization running in native mode.

NDR
See definition for: non-delivery report

Network News Transfer Protocol
(NNTP) An application protocol used in TCP/IP networks. Enables clients to read and post information to USENET newsgroups.

See also: newsgroup

news site
A collection of related newsgroups.

newsfeed
The flow of items from one USENET site to another.

newsgroup
An Internet discussion group that focuses on a particular category of interest.

See also: Network News Transfer Protocol

NNTP
See definition for: Network News Transfer Protocol

non-delivery report
(NDR) A notice that a message was not delivered to the recipient.

NTFS file system
The file system designed for use specifically with the Windows NT operating system. NTFS supports file system recovery and extremely large storage media. It also supports object-oriented applications by treating all files as objects with user-defined and system-defined attributes.

NTLM authentication protocol
A challenge/response authentication protocol. The NTLM authentication protocol was the default for network authentication in Windows NT version 4.0 and earlier. The protocol continues to be supported in Windows 2000 but is no longer the default.

O

object
The basic unit of Active Directory. It is a distinct, named set of attributes that represents something concrete, such as a user, a printer, a computer, or an application.

one-step migration

One of two migration methods available in Migration Wizard. In a one-step migration, Migration Wizard extracts migration files from another messaging system server and then imports the migration files to Exchange in one operation.

See also: two-step migration

operations master

A domain controller that has been assigned one or more special roles in an Active Directory domain. The domain controllers assigned these roles perform operations that are single-master (not permitted to occur at different places in the network at the same time). Examples of these operations include resource identifier allocation, schema modification, primary domain controller election and certain infrastructure changes. The domain controller that controls the particular operation owns the operations master role for that operation. The ownership of these operations master roles can be transferred to other domain controllers.

organization

A set of computers running Microsoft Exchange Server that provide messaging and collaboration services within a business, an association, or a group.

organizational unit

An Active Directory container into which you can place objects such as user accounts, groups, computers, printers, applications, file shares, and other organizational units. Organizational units can be used to contain and assign specific permissions to groups of objects, such as users and printers. An organizational unit cannot contain objects from other domains. An organizational unit is the smallest unit you can assign or delegate administrative authority to.

Outlook Web Access

Outlook Web Access for Microsoft Exchange 2000 Server provides users access to e-mail, personal calendars, group scheduling, contacts, and collaboration applications using a Web browser. It can be used for UNIX and Macintosh users, users without access to a Microsoft Outlook 2000 client, or users connecting from the Internet. Outlook Web Access offers cross-platform client access for roaming users, users with limited hardware resources, and users who do not have access to their own computers.

P

permission

Authorization for a user to perform an action, such as sending e-mail for another user or posting items in a public folder.

PKI

See definition ***for:*** public key infrastructure

Point-to-Point Tunneling Protocol
(PPTP) An encryption protocol used for remote computers to securely access other computer networks across an Internet connection. Often used with Virtual Private Networks (VPNs).

policy
A collection of configuration settings that are applied to one or more Exchange configuration objects. You can use policies to simplify the administration of Exchange. You can define a policy that controls the configuration of some or all settings across a server or other objects in an Exchange organization. After policies are defined and implemented, editing the policy and applying the changes will change the configuration of all servers and objects covered by the policy.

POP3
See definition for: Post Office Protocol version 3

Post Office Protocol version 3
(POP3) An Internet protocol that allows a client to download mail from an inbox on a server to the client computer where messages are managed. This protocol works well for computers that are unable to maintain a continuous connection to a server.

PPTP
See definition for: Point-to-Point Tunneling Protocol

presence information
In Instant Messaging, the information visible to users that shows the online status of contacts (online, busy, away, and so on).

primary connection agreement
A connection agreement that matches objects that exist, and creates new objects that did not exist.

See also: connection agreement

primary domain controller emulator master
The domain controller assigned to act as a Microsoft Windows NT primary domain controller (also known as PDC) to service network clients that do not have Active Directory client software installed, and to replicate directory changes to any Windows NT backup domain controllers (also known as BDCs) in the domain. For a Windows 2000 domain operating in native mode, the primary domain controller emulator master receives preferential replication of password changes performed by other domain controllers in the domain and handles any password authentication requests that fail at the local domain controller. At any time, there can be only one primary domain controller emulator in a particular domain.

proxy server
A firewall component that manages Internet traffic to and from a LAN and can provide other features, such as document caching and access control.

public folder
A folder that coworkers can use to share a wide range of information, such as project and work information, discussions about a general subject, and classified ads. Access permissions determine who can view and use the folder. Public folders are stored on computers running Exchange.

public folder hierarchy
A tree or hierarchy of public folders with a single public folder store.

public folder replication
The process of keeping copies of public folders on other servers up to date and synchronized with each other.

public folder store
The part of Microsoft Web Storage System that maintains information in public folders. A public folder store consists of a rich-text .edb file, plus a streaming native Internet content .stm file.

public key infrastructure
(PKI) The laws, policies, standards, and software that regulate or manipulate certificates and public and private keys. In practice, it is a system of digital certificates, certification authorities, and other registration authorities that verify the validity of each party involved in an electronic transaction. Exchange Key Management Service (KMS) works with Windows 2000 Certificate Services to provide a PKI for Exchange organizations. Through a third-party certification authority, a Windows 2000 Certificate Services server may be part of a larger PKI that extends beyond an organization. Certificate Services issues X.509 version 3 certificates that bind a user s identity, such as an e-mail address and distinguished name, to their public keys. KMS maintains an encrypted database of the corresponding private encryption keys.

See also: certificate template, key pair

R

RAID
See definition for: redundant array of independent disks

recipient
An Active Directory object that is mail-enabled, mailbox-enabled, or that can receive e-mail. A recipient is an object within Active Directory that can take advantage of Exchange functionality.

recipient policy
Policies that are applied to mail-enabled objects to generate e-mail addresses. They can be defined to apply to thousands of users, groups, and contacts in Active Directory by using a Lightweight Directory Access Protocol (LDAP) query interface in a single operation.

Recipient Update Service
An Exchange 2000 service that updates the recipient objects within a domain with specific types of information. You can schedule appropriate intervals to update the recipient objects. For example, this service updates recipient objects with address list membership and e-mail addresses at intervals scheduled by the administrator.

redundant array of independent disks
(RAID) A mechanism for storing identical data on multiple disks for redundancy, improved performance, and increased mean time between failures (MTBF). RAID provides fault tolerance and appears to the operating system as a single logical drive.

remote bridgehead server
A server that handles e-mail flow to and from a routing group connector in a different routing group.

remote procedure call
(RPC) A routine that transfers functions and data among computers on a network.

replica
A copy of a public folder that contains all of the folder's contents, permissions, and design elements, such as forms behavior and views. Replicas are useful for distributing user load on servers, distributing public folders geographically, and for backing up public folder data.

replication
See definition for: directory replication

reverse proxy server
A reverse proxy server is similar to a regular proxy server used for outbound network traffic except that it relays connection requests for inbound network traffic.

routing group
A collection of Exchange servers that have full-time, reliable connections. Messages sent between any two servers within a routing group go directly from source to destination. Similar to administrative groups, routing groups are optional and are not visible in System Manager unless they are enabled.

routing group bridgehead server
A server within a routing group that exchanges directory updates with a server in another routing group.

routing group connector
A connector that specifies the connection of a local routing group to a server in a remote routing group. It also specifies the local bridgehead server, if any, and the connection cost, schedule, and other configuration properties.

RPC
See definition for: remote procedure call

RSA cryptographic algorithms
A widely used set of public key algorithms that are available from RSA Data Security, Inc. The RSA cryptographic algorithms are supported by the Microsoft Base Cryptographic Service Provider and the Microsoft Enhanced Cryptographic Service Provider.

S

schema
A logical model for data; an organizational framework. Schema defines the universe of objects that can be stored in Active Directory. For each object class, the schema defines what attributes an instance of the class must have, what additional attributes it can have, and what object class can be a parent of the current object class.

schema master
The domain controller that performs write operations to the directory schema. Schema updates are replicated from the schema master to all other domain controllers in the forest. Only the schema master domain controller can perform this task.

Secure Sockets Layer
(SSL) A protocol designed to establish a secure communications channel to prevent the interception of critical information, such as credit card numbers.

security association
A set of parameters that defines the services and mechanisms necessary to protect Internet Protocol security (IPSec) communications.

security descriptor
In Windows 2000, it is possible to set security for objects because every object has a security descriptor. The security descriptor is where the security settings for the object are stored. A security descriptor consists of the security identifier (SID) of the object owner, a group SID used by the Portable Operating System Interface (POSIX) subsystem and Services for Macintosh, a discretionary access control list (DACL), and a system access control list (SACL).

security identifier
(SID) A data structure of variable length that uniquely identifies user, group, service, and computer accounts within a forest. Every account is issued a SID when the account is first created. Access control mechanisms in Windows 2000 identify security principals by SID rather than by name.

security subsystem
See definition *for:* Local Security Authority

server cluster
A group of independent computers that work together to run a common set of applications. The computers are physically connected by cables and programmatically connected by cluster software. These connections allow the computers to use problem-solving features, such as load balancing, while appearing to the user and applications as a single system.

SID
See definition *for:* security identifier

Simple Mail Transfer Protocol
(SMTP) The standard protocol for Internet mail. SMTP transfers mail from server to server and from mail system to mail system. In Exchange 2000 Server, SMTP is the native transport protocol.

site
In Windows 2000, one or more reliable and fast TCP/IP subnets. Setting up Windows 2000 sites allows you to configure Active Directory access and a replication topology to take advantage of the physical network.

site (in earlier versions of Exchange)
A group of servers (usually in the same geographic location) that share the same directory information and can communicate over high-bandwidth, permanent, and synchronous connections.

site link
An Active Directory object that represents a set of sites that can communicate at uniform cost. For Internet Protocol (IP) transport, a typical site link connects just two sites and corresponds to an actual WAN link. An IP site link connecting more than two sites might correspond to an ATM backbone connecting more than two clusters of buildings on a large campus, or several offices in a large metropolitan area connected by leased lines and IP routers.

Site Replication Service
A directory service (similar to the directory used in Exchange Server 5.5) implemented in Exchange 2000 to allow integration with Exchange 5.x sites that use both remote procedure call (RPC) and mail-based replication. Site Replication Service works with Active Directory Connector (ADC) to provide replication services from Active Directory to the Exchange 5.x Directory Service.

smart host
A designated server through which Exchange routes all outgoing messages. The smart host then makes the remote connection. If a smart host is designated, the Exchange server only needs to transmit to the smart host, instead of repeatedly contacting the domain until a connection is made. Also known as a relay host.

snap-in
Software that makes up the smallest unit of a Microsoft Management Console (MMC) extension. One snap-in represents one unit of management behavior.

SSL
See definition for: Secure Sockets Layer

SSPI
See definition for: Microsoft Security Service Provider Interface

storage group
A collection of mailbox stores and public folder stores that share a set of transaction log files. Exchange manages each storage group with a separate server process.

system policies
Policies that apply to server-side objects, such as mailbox stores, public folder stores, and servers.

T

TCP/IP filtering
A feature of Windows 2000 TCP/IP that allows you to specify exactly which types of incoming non-transit IP traffic are processed for each IP interface.

TLS
See definition for: Transport Layer Security

token
A random character string given to users to enable advanced security for them.

transaction log file
A file that maintains a record of every message stored in a storage group and provides fault tolerance in the event that a database must be restored.

transitive trust relationship
The trust relationship that inherently exists between Windows 2000 domains in a domain tree or forest, or between trees in a forest, or between forests. When a domain joins an existing forest or domain tree, a transitive trust is automatically established. Windows 2000 transitive trusts are always two-way relationships.

Transport Layer Security
(TLS) A generic encryption technology similar to Secure Sockets Layering (SSL). Like SSL, TLS encrypts information over the wire between a client and a server to prevent packet-sniffing and other attempted security breaches. In Exchange, TLS is used by SMTP virtual servers. While authentication prevents the server from unauthorized access, TLS encryption protects the information it sends and receives.

tree

Also known as a directory hierarchy. A hierarchical arrangement of one or more Windows 2000 domains that share a common naming structure. End points on the tree are usually objects. Nodes in the tree, or points at which the tree branches, are containers that hold a group of objects or other containers. A tree shows how objects are related to one another.

trust relationship

The relationship between two domains that makes it possible for a user in one domain to access resources in another domain.

two-step migration

One of two migration methods available in Migration Wizard. In a two-step migration, first you extract migration files from another messaging system server and, if necessary, review or edit the migration files. Then you import the migration files to Exchange.

***See** also:* one-step migration

universal group

A Windows 2000 group available only in native mode that is valid anywhere in a forest. A universal group appears in the global catalog but contains primarily global groups from domains in a forest. This is the simplest form of group and can contain other universal groups, global groups, and users.

***See** also:* domain local group

user

An Active Directory object that has a Windows security account and a password. A user is the only Active Directory object that can have a mailbox associated with it. A user in Windows 2000 is the equivalent of a mailbox in earlier versions of Exchange.

V

virtual root

A mapping between a specific path or name and a physical storage location, be it a local file directory network share or redirection to another URL. For Hypertext Transfer Protocol (HTTP), a virtual root defines a mapping between a URL path and a physical storage location. For Network News Transfer Protocol (NNTP), a virtual root defines a mapping between a news group name and a physical storage location.

virtual server

A collection of services that appears to clients as a physical server. It is an instance of a protocol service (for example, SMTP) with a defined set of Internet Protocol (IP) address/port combinations and an independent collection of configuration properties. A virtual server typically includes all the resources necessary to run a particular application, including a network name resource and an IP address resource.

Web Distributed Authoring and Versioning
(WebDAV) An extension of Hypertext Transfer Protocol (HTTP) 1.1 that allows clients to perform remote Web content authoring. Content that is stored on a server can be accessed by a client through HTTP by using WebDAV extensions. The client can perform tasks provided by HTTP, including reading e-mail and documents. If the client also supports WebDAV, the client can manipulate mail, change calendar appointments, modify and create new documents on the Exchange 2000 server, and create Web-based forms. WebDAV uses Extensible Markup Language (XML) as the format for transmitting data elements.

Web Storage System
A storage platform that provides a single repository for managing multiple types of unstructured information within one infrastructure. Microsoft Web Storage System combines the features and functionality of the file system, the Web, and a collaboration server (such as Exchange Server) through a single, URL-addressable location for storing, accessing, and managing information, as well as building and running applications.

WebDAV
See definition for: Web Distributed Authoring and Versioning

X

X.400 Connector
A Microsoft Exchange Server component that is integrated with the message transfer agent (MTA) and can be configured to connect routing groups within Exchange, or to route messages to other X.400 systems. When handling communication between Exchange and other X.400 systems, it maps addresses and converts Exchange messages to native X.400 messages and vice versa.

Z

zone
In a DNS database, a zone is a contiguous portion of the DNS tree that is administered as a single separate entity by a DNS server. The zone contains resource records for all the names within the zone.

Index

access control
  DACLs (discretionary access control lists) 829
  how access control works 836–838
  model 838
  public folders 293–294
  security descriptors 836
ACEs (access control entries) 86, 837–838
ACLs (access control lists) 86,259–260
Active Directory
  access to data 893–894
  Active Directory Connector ***See*** ADC
  Active Directory Service Interface ***See*** ADSI
  auditing objects 839
  backbone scenarios 496
  backup prerequisites 769
  backups, performing 770
  bulk changes 895
  Certificate Services 842
  Cleanup Wizard 154–155
  client access 124–127
  coexistence ***See*** coexistence
  compared to Exchange 5.5 Directory Service 111
  compared to Security Accounts Manager (SAM) 829
  configuration partitions 897
  connection agreements 130–134
  connectivity 407–409
  DACLs (discretionary access control lists) 829
  data partitions 895–897
  Delegation of Control Wizard 285
Active Directory *(continued)*
  dependencies 706
  deployment environment 47–49
  deployment tips 148
  detecting changes 704
  directory partitions 110
  disaster recovery 793
  distribution groups 101, 229, 293, 831–833
  domain controllers ***See*** domain controllers
  Domain Name System ***See*** DNS
  DomainPrep ***See*** DomainPrep
  domains ***See*** domains
  ForestPrep ***See*** ForestPrep
  forests ***See*** forests
  GAL (global address list) 120
  global catalog servers ***See*** global catalog servers
  groups 100–107,831–836
  groups vs. distribution lists 120
  initial replication 703
  Instant Messaging 934–935
  logical components 96
  logical structure 71, 91
  LSA (Local Security Authority) 830–831
  management agent for Metadirectory Services 457
  message categorizer 387
  migration scenarios 154–157
  Migration Tool 155–157
  multi-master replication 110
  namespaces ***See*** namespaces

Active Directory *(continued)*
naming contexts 96–98
object class mappings 704–705
object definitions 407–408
objects that replicate 705
optimizing 871–872
organizational units ***See*** organizational units
Outlook 2000 access to data 893
Outlook 98 and Outlook 97 access to data 894
overview 63–64, 109
partial replicas 110
partitions for hosting other companies 507–508
permissions 469
physical structure 91
POP3 user mailboxes 713
populating user accounts 153–157
preparing for mixed mode 142–144
public folders 398–399. 681, 922
recipient policies 408
Recipient Update Service 123,409
replicas 110
replication 110–111
requirements 248–251
schema ***See*** schema, Active Directory
security groups 101,229,293, 831–833
server directory access 894
services, list of 120
specifying servers 384–385
storing data 895–898
trees ***See*** trees
troubleshooting 953–954
trust relationships ***See*** trust relationships
Users and Computers 294–295,953
what's new 4–5, 702
Active Directory Cleanup Wizard 154–155
Active Directory Connector ***See*** ADC
Active Directory Delegation of Control Wizard 285
Active Directory Migration Tool 155–157
Active Directory Service Interface ***See*** ADSI
Active Directory Users and Computers
moving mailboxes 294–295
troubleshooting 953
Active Server Pages (ASP)
overview 732
sample 732–733
what's new 16
active/active clustering 12,679, 764–766
ActiveX Data Objects ***See*** ADO
ADC (Active Directory Connector)
accounts preparation 141
Active Directory Cleanup Wizard 154–155
Active Directory Migration Tool 155–157
architecture 702–703
backbone scenarios 496–498
configuration connection agreements 133
connection agreement credentials 864
connection agreement operations 134–135
consolidation of mailboxes 129
contingency plans 142
deployment environment 49
deployment tips 148
enhanced 128
event IDs 146–147
features 127
in-place domain upgrades 154
installation locations 138
installation permissions 863
installation requirements 145
inter-organization synchronization 452-453
large company deployment 435

ADC (Active Directory Connector) *(continued)*
logged events 146
migrating accounts to Windows 2000 154–157
mixed mode administration 298
mixed mode connectivity 425
monitoring 146–147
moving custom recipients to Active Directory 257
multiple instances 138
object matching 135–137
overview 127
permissions 863–864
preparing for deployment 140–144, 152–153
questions before deployment 139–140
reasons to install 130
replicating distribution lists to groups 106
replicating files from Exchange 5.5 703–705
replication overview 137
server preparation 141
service account permissions 863
sIDHistory attribute 156–157
synchronizing Active Directory with Exchange 5.5 451
testing migration paths 191
troubleshooting 954–955
user connection agreements 131–132
versions of 128,453
viewing connection agreements 130
Windows 2000 128
additional resources *See* Web sites
address hooks
compared to address list views 258
default views 400
deployment environment 51
offline 255
searches 119
address books *(continued)*
views in earlier versions 121
address lists
compared to address books views 258
compatibility with earlier versions 122, 401
default 122
default address book views 400
global *See* GAL (global address list)
offline 123
overview 121,400
Recipient Update Service 401
testing 188
addresses
deployment environment 230
Lotus Notes 228
resolving IP 911–912
types of 230
user 934
administration
Active Directory Delegation of Control Wizard 285
ADC (Active Directory Connector) 298
administrative groups 297
administrative models *See* administrative models
administrative tools 284–285
delegation of server 5
distribution groups 293
Event Viewer 303–304
Exchange Administration Delegation Wizard 286, 292
Exchange Monitoring and Status Tool 300–302
group responsibilities 288
hosting other companies 516–517
levels 858–859
mail-enabled groups vs. distribution lists 293

administration *(continued)*
Message Tracking Center 302
mixed mode 297
monitoring performance 299–304
moving Exchange 5.5 servers 298–299
moving mailboxes 294–297
optimizing 898–899
organizational units 282–283
overview 281
Performance Logs and Alerts 303
permissions 286
public folders 261, 293–294, 298
Queue Viewer 302
reactive vs. proactive tools 300
recipient management 282
recipient policies 287
roles 858
security groups 293
server management features 290–292
server management overview 287
server responsibilities 287–288
System Monitor 303
user 5
what's new 5–6, 281
Windows 2000 tools 302–304
administrative groups
branch office deployment 464
centralized administrative model 472
compared to sites 375
distributed administrative model 474
example of 464465
mixed administrative model 477
mixed mode 291
naming conventions 226
server management 290
administrative groups *(continued)*
testing 189
what's new 6
administrative models
branch office centralized model 469–473
branch office distributed model 473475
branch office mixed model 475–477
branch office security 468
branch office, list of 467
branch office, planning for 467469
centralized administration 283
centralized management 289
delegated administration 284
distributed management 289
messaging group administration of attributes 284
mixed mode 290
recipient management overview 283
types of server models 288
what's new 3–9
administrative requirements 42
ADO (ActiveX Data Objects)
accessing objects through ADSI 718–719
navigating through Web Storage System 714
object collection 715–716
querying Web Storage System 716–717
what's new 16
when to use 726
ADSI (Active Directory Service Interface)
accessing ADO objects 718–719
changing items 721
creating temporary objects 721
directory structure 720
enumerating objects in containers 722
error messages 723
object instance and schema definitions 721

ADSI (Active Directory Service Interface) *(continued)*
object properties 721
overview 719–720
retrieving information 721–723
schema scripts 882–883
SQL object definitions 721
advanced queuing engine 387,491, 669
affinity, public folder
administration 294
client connection method 682
connectors 59–60
deployment environment 231–232
tuning 399
upgrading 262
alerts, virus 332
algorithms
encryption 853
link state *See* link state algorithms
anonymous access 568
antivirus solutions
See also virus protection
adding scanning to monitoring services 332
checklists 329
connector servers 367
performance counters for scanning 331
testing 329–331
vendors 329–332
API components
ADO 714–719
ADSI 719–723
CDO 723–726
list of 709.714
MAPI 727
OLE DB 727
overview 714
Application Deployment Wizard 757
application development
event support 738–749
overview 707
tools 756–757
Web-based collaboration applications 731–737
Web-based collaboration overview 708–730
workflow components 750–756
arbitration 169–170
architecture
ADC (Active Directory Connector) 702–703
back-end *See* back-end architecture
connector 363
deployment environment 49–55, 210
Exchange Server 5.5 700–701
ExIPC 661–664
front-end *See* front-end architecture
high availability 333–334
hosting 505–506
IFS (Installable File System) 677–679
IIS process 661
IMAP4 672–673
Lotus Notes and Exchange coexistence 236–238
message flow 683
message routing 683–684
message transport 684
Metadirectory Services 456–457
mixed mode 700
namespaces 66
NNTP 673–674
Outlook Web Access 557–558,693
overview 659
POP3 672–673
previous versions of Exchange 700
public folders 679

architecture (continued)
review 49–55
scalable and resilient 347
servers 333–334, 347,659–660
SMTP 665
Web Storage System 661,674
WebDAV 698–699
what's new 700–701
archiving backups 801
ASP (Active Server Pages)
overview 732
sample 732–733
what's new 16
assessing project risks 20–22
asymmetric cipher 852
attachments with viruses 330
auditing 493–495, 838–839
authentication
Basic 568
delegation of 840–841
dual 569
Integrated Windows 568
Kerberos 839–841
Outlook Web Access 567–569
pass-through 569
shortcut trust relationships 96
availability, server
architecture 333–334, 347
creating processes 348
implementing 339–347
overview 333
pilot programs 196
requirements 334–338
summary 349

B

backbone configurations
address lists 400–401
diagrams 53
directory access 384–387
overview 383
public folders 396–399
routing and transport 387–396
scenarios See backbone scenarios
virus protection 323–324
backbone scenarios
Active Directory 496
ADC (Active Directory Connector) 496–498
advanced queuing engine 491
advanced SMTP command verbs 488–489
auditing your system 493–495
Chat Service 492
chunking 488–489
connectors to other mail systems 503
Data Conferencing Provider 492–493
directory synchronization 494
directory synchronization between forests 499
directory topology 496–499
Financial Bank, Inc. overview 487–488
global catalog servers 498–499
Instant Messaging 492
Internet mail architecture 494–495
link state algorithms 491, 501
message categorizer 491
message routing 490–491,494,503
messaging coexistence 503
messaging topology 500–503
name resolution 499
overview 487

backbone scenarios *(continued)*
pipelining 489
public folders 490
routing groups 500–503
storage in native Internet format 489
summary 504
Web Storage System 489–490
back-end architecture
centralized administrative model 470–471
configuring Outlook Web Access servers 565
implementing availability technologies 345–347
mixed administrative model 475476
Outlook Web Access dual authentication 569
Outlook Web Access firewalls 12, 563
Outlook Web Access overview 562, 565, 697–698
Outlook Web Access pass-through authentication 569
scaling servers 886
backup and restore
backup categories 770–772
backup compared to Exchange 5.5 307
backup prerequisites 766–770
backup process 772–777
backup to disk 309–310
backup types 314–315
best practices 800–802
branch office servers 311–312
centralized servers 315–317
data center installations 315–317
database consistency 784–791
database technology 760–766
databases vs. storage groups 306–307
deployment environment 230
design issues 309–311
backup and restore *(continued)*
design scenarios 311–317
disaster recovery 791–800
disaster recovery prerequisites 779–780
Exchange 5.5 vs. Exchange 2000 305–309
failover 761
large departmental servers 312–315
overview 759
parallel operations 308
recommendations 309
restore categories 781–783
restore compared to Exchange 5.5 307–308
restore constraints 783–784
restore described 777–778
server design overview 305
server recovery 309
small organization servers 311–312
testing 187
throughput 310–311
troubleshooting 970–971
bandwidth
***See** also* traffic analysis
low-bandwidth environments 395–396.447
network bandwidth for mailboxes 295
Basic authentication 568
best practices
backup and restore 800–802
connectivity 403–407
full-text indexing 890–892
block cipher 853
boot sector viruses 322
bottlenecks, monitoring 823–825
branch office deployment
administrative groups 464
administrative model security 468

branch office deployment *(continued)*
administrative models, list of 467
backup and restore 311–312
centralized administrative model 469–473
client comparisons 484–486
client considerations 478
distributed administrative model 473475
functionality requirements 479
global catalog server placements 477478
IMAP4 481–482
Instant Messaging 447
mixed administrative model 475–477
mobile vs. offline requirements 479–486
objectives 479
Outlook 2000 479481
Outlook Web Access 483
overview 463
planning for administrative models 467-469
POP3 481–482
replication 463
routing groups 464467
Terminal Services 484486
Windows 2000 dependencies 463467
bridgehead servers
deployment environment 52–53
link state propagation 915
multiple 911
recovering connections 915
resolving IP addresses 911–912
routing groups 466
security 911
SMTP connector 912
building project plans
deploying phase 34–35
developing phase 32–34
building project plans *(continued)*
envisioning phase 22–27
overview 19
phased approach 20–22
planning phase 27–31
summary 35
business requirements 4143

C

CA (certification authority) 524–525
cache
directory 386
disk controller cache settings 359–360, 367
DSAccess (Directory Service Access) 900–903
Instant Messaging 932
parameters 386–387
shared 384
Time to Live setting (CacheTTL) 901
calendaring 416, 615–618
canonical order 838
cc:Mail *See* Lotus cc:Mail
CD, Resource Kit companion ix
CDO (Collaboration Data Objects)
ADO integration 724
collaboration-specific data 724
contact information 725
IDataSource interface 724
object and programming models 724
overview 723
scheduling 725–726
sending mail 726
what's new 14–15
when to use 726
workflow event sinks 751–752
Workflow Objects 15

centralized administration model 283
centralized branch office administrative model 469–473
centralized management model 289
certificate revocation list (CRL)
high and very high trust levels 545
publishing to a trusted third-party 554–556
Certificate Services
Key Management Service 853–854
security 842
certificate trust list (CTL) 554
certificates
commercial vs. self-hosted root CA 534
enrollments 524,527–530,547–548
overview 523–525
renewals 530
revocations 525.531–532
certification authority (CA) 524–525
chat servers
capacity 355
configurations 944
described 352
disk configurations 369
locations 434-436
memory 372
Chat Service
backbone scenarios 492
client connections 945
deploying 432433
firewalls 436
network address translations 436
overview 431432,944
recommendations 437
scalability 434, 945
security 436
Chat Service *(continued)*
server configurations 944
server locations 434–436
TCP ports 436
upgrading 275
what's new 12
Chatmig.exe 275
checkpoint files 763
checksum 764
chunking 488489,671–672
ciphers 852–853
circular logging 262,763
Cleanup Wizard 154–155
clients
Active Directory 124–127
branch office deployment 478
deployment environment 233–234
directory access 4
mail access at branch offices 468
network traffic ***See*** traffic analysis
Outlook Web Access 561–562
server access 58
server communication security 865–866
software test plans 183–185
support 11
troubleshooting public folders 959
upgrading 276–277
clustering
active/active 12,679, 764–766
cluster server node backups 772
component support 765–766
database technology 764–766
deployment environment 56
failover 765
high availability for servers 344–345

clustering *(continued)*
recovering cluster servers 797–800
recovering lost cluster quorums 799–800
what's new 701
coexistence
See also Exchange 5.5
Active Directory Connector *See* ADC
administrative groups 291
administrative tools 899
architecture for Exchange Server 5.5 700–701
architecture for Lotus Notes and Exchange 236–238
architecture for mixed mode 700
connection agreements *See* connection agreements
Lotus Notes and Exchange overview 234–235
mailbox upgrades 163–164
matching objects 135–137
messaging in backbone scenarios 503
migration process for Lotus Notes and Exchange 240
routing groups 291
services for Lotus Notes and Exchange 235
topology for Lotus Notes and Exchange 237
collaboration applications *See* Web-based collaboration
Collaboration Data Objects *See* CDO
commercial vs. self-hosted root CA 534
company structure 40–41
Conference Management Service 13
Conferencing Server
Data Conferencing Provider 945
overview 945
what's new 13–14
configuration connection agreements 133
configuration data for backups 770
configuration naming contexts 97
configurations
backbone *See* backbone configurations
Chat Service servers 944
DSAccess (Directory Service Access) 903–904
Lotus cc:Mail 420422
Lotus Notes 411414
Microsoft Mail 416–417
Novell GroupWise 423424
configuring
administrator computers 899
back-end servers 565
diagnostic logging 146
DNS (Domain Name System) 936–938
front-end servers 565–567
Instant Messaging 934–935
Instant Messaging server components 938–939
Outlook Web Access servers 565–567
public folders 398
unified namespaces 82–83
connecting to public folders 681–683
connection agreements
configuration 133
creating initial 134
operations 134–135
overview 130
types of 130–131
user 131–132
connectivity
See also connectors
Active Directory 407–409
best practices 403–407
connecting Lotus Notes and Exchange 405
Internet 213–214,219
Metadirectory Services 427

connectivity *(continued)*
mixed mode 424–425
overview 403
planning 403–407
PROFS 426–427
retail stores 213–214
SNADS 426–427
troubleshooting 960–968
what's new 9
connector servers
antivirus solutions 367
capacity 354–355
deployment environment 53. 224
described 352
disk configurations 362–367
disk controller cache settings 367
location 366–367
memory 370
network configurations 372
RAID 367
specifications 224
connectors
See also connectivity
architecture 363
Exchange 2000 375
Exchange 2000 vs. Exchange 5.5 388–389
list of 362, 409
migrating 425
other e-mail systems 9, 503
performance 395
permissions 864
public folder affinity 60
routing groups 906
types 388–389
upgrading 265–275
contacts vs. custom recipients 257
conventions, document x
copy backups 771
corporate backbones *See* backbone configurations; backbone scenarios
costs
additional domains 73–74
budgeting for test labs 178
downtime 335–336
link state tables 913
multiple routing groups 907,917–918
pilot programs 195
public folders 682
test lab hardware 179
virus protection 319
counters
data-link layer 818
full-text indexing 892
logical hard disk 814
monitoring large numbers 811
network layer 818
physical hard disk 814
presentation/program layer 819
recommended thresholds 823
transport layer 818–819
virus scanning 331
creating availability processes 348
CRL (certificate revocation list)
high and very high trust levels 545
publishing to a trusted third-party 554–556
CryptoAPI 526
cryptographic service providers 526
CTL (certificate trust list) 554
custom recipients vs. contacts 257

D

DACLs (discretionary access control lists) 829
daily backups 767
data
  configuration 770
  dynamic 768
  restoring data to non-production servers 784
  static 767
  types of data to back up 767
  verifying backups 801
Data Conferencing Provider
  backbone scenarios 492–493
  Conferencing Server 945
  deployment environment 220–221
  Internet connectivity 220
  number of servers 221
  security 220
  server locations 220
  what's new 13
  Windows 2000 Server dependencies 945
data conferencing servers
  capacity 355
  deployment environment 225
  described 352
  disk configurations 368
  locations 220
  memory 371
  number of 221
  specifications 225
data storage
  multiple databases 340–341, 675–676
  what's new 10–11
database technology
  checkpoint files 763
database technology *(continued)*
  checksum 764
  circular logging 763
  database files 761–762
  database GUID 762
  Extensible Storage Engine (ESE) 760
  failover 761
  individual database backups 764
  log file signatures 763
  mailbox GUID 762
  message hash 764
  multiple storage groups 761
  overview 760
  server clusters 764–766
  transaction log files 762
databases
  basic schema 789
  checking integrity 788
  compared to storage groups 306–307
  consistency 784–791
  defragmenting 785–786
  multiple ***See*** multiple databases
  recovering in Web Storage System 795
  reference 790
  repairing 788
  restoring Key Management Service 797
  restoring single 783
  restoring SRS 796
  sizing for global catalog servers 112–115
  soft recovery 786,789
  transaction log files 786
  upgrade benefits vs. risks 163
data-link layer in Network Monitor 818
decryption ***See*** encryption
default address lists 122

defragmenting databases 785–786
delegated administration model 284
delegation of authentication 840–841
Delegation of Control Wizard 285
deploying
  branch offices ***See*** branch office deployment
  Chat Service 432–433
  deploying phase described 20
  forests 89
  Instant Messaging 438–441
  projects 34–35
deployment environment
  See *also* deployment strategies
  Active Directory 47–49
  administration 230
  architectural review 49–55, 210
  architecture for Lotus Notes and Exchange 236–238
  assessment phases 202
  backup and restore 230
  branch office ***See*** branch office deployment
  business requirements 41–43
  client access 233–234
  company structure 40–41
  connector servers 224
  Data Conferencing Provider 220–221
  data conferencing servers 225
  design goals 203
  design goals for clients 233
  design plans 205–207
  design summary 242–243
  DNS (Domain Name System) 46
  domain controllers 209
  domain design 208–209
  end-user input 209
deployment environment (continued)
  forest design 208
  front-end servers 213–214.224–225
  gap analysis 203–205
  global catalog considerations 208–209
  global catalog placement 214
  IMAP4 234
  Internet connectivity 213–214, 219
  LAN 45
  LitWare example overview 201
  Lotus Notes and Exchange migration overview 234–235
  mailbox servers 221–223
  message routing 214–218
  migrating Lotus Notes to Exchange 239–240
  multiple forests 171–174
  naming conventions 209,225–230
  NetMeeting 234
  network infrastructure 43
  organizational forms 61
  Outlook 2000 234
  Outlook Web Access 234
  overview 39
  physical networks 207–208
  POP3 234
  preparing the environment 207–209
  public folder servers 223
  public folders 59–61, 231–233
  real-time collaboration services 219–221
  required functionality 202–203
  retail store connectivity 213–214
  server configurations 55–58
  server locations 210–212
  server roles 221–225
  services for Lotus Notes and Exchange 235

deployment environment *(continued)*
summary 62
training administrators and users 241
WAN 44
Windows NT 4.0 vs. Windows 2000 46 47
deployment planning templates 22
deployment scenarios
choosing strategies 171
Instant Messaging 940–944
inter-company e-mail security 550–552
inter-company e-mail security with trusted third party 552–556
intra-company e-mail security 546–550
message routing 214–218
multiple organizations 174
routing groups 501–502,909–910
separate forest 172–174
separate organizations in forest 171–172
server locations 210–212
deployment strategies
See also deployment environment
ADC preparations 140–144
ADC questions 139–140
ADC tips 148
existing organization 150–151
multiple organizations 174
new organization 150
Outlook Web Access 562–564
overview 149
populating user accounts in Active Directory 153–157
preparing to deploy 151–153
roadmap 149–151
scenarios *See* deployment scenarios
separate forest 172–174
deployment strategies *(continued)*
separate organizations in forest 171–172
upgrading mailboxes 161–171
upgrading Windows NT servers 157–161
design considerations
backup and restore overview 305
backup and restore throughput 310–311
backup to disk 309–310
deployment environment 50
routing groups 908
server upgrades 159
sites 51
validating in project planning 32
what's new 5
design scenarios, backup and restore
branch office servers 311–312
centralized servers 315–317
data center installations 315–317
large departmental servers 312–315
small organization servers 311–312
design scenarios, S/MIME
inter-company 537–540
inter-company with trusted third party 541–545
intra-company 532–536
list of 532
desktop virus protection 324
developing phase
building systems 33
described 20
pre-pilot tests 33
project overview 32
user pilot tests 33–34
validating designs 32
diagnostic logging 146, 806–808
dial-up access 251

differential backups 771
Digital Dashboard 233
digital signatures 523
directories
appearance 5
SMTP 666–667
topology for backbone scenarios 496–499
directory cache 386
Directory Service Access ***See*** DSAccess
directory synchronization
backbone scenarios 494,499
Exchange 2000 inter-organization solutions 453–454
Exchange 5.5 inter-organization solutions 450–453
inter-organization overview 449
inter-organization replication vs. synchronization 450
Lotus cc:Mail 418–419
Lotus Notes 410-411
Microsoft Mail 414–416
Novell GroupWise 422
PROFS 427
SNADS 427
disaster recovery
Active Directory 793
cluster servers 797–800
Key Management Service 796–797
lost cluster quorums 799–800
member servers 793–797
overview 791–792
performance factors 780
planning 54–55
prerequisites 779–780
requirements 792
server recovery 309
disaster recovery ***(continued)***
Site Replication Service (SRS) 795–796
discretionary access control lists (DACLs) 829
disk configurations
chat servers 369
connector servers 362–367
data conferencing servers 368
front-end servers 367
Instant Messaging servers 368–369
mailbox servers 357–362
overview 356–357
public folder servers 357–362
video conferencing servers 368
distributed branch office administrative model 473475
distributed management model 289
distributed services 339
distribution groups
advantages 101
described 229
expansion 388
scope 832–833
distribution lists
compared to groups 120
compared to mail-enabled groups 5, 293
migrating to Active Directory 257
testing 190
upgrading 170–171,260
DNS (Domain Name System)
See also namespaces
advantages of DNS service 76-77
configuring 936–938
configuring unified namespaces 82–83
deployment environment 46
dynamic updates 77
namespace overview 70

DNS (Domain Name System) *(continued)*
naming recommendations 79–82
non-Windows DNS service 77–78
placing DNS servers 78
planning namespaces 65–66
record example 83
resource records 70
security 77,868–869
service for clients 76
service requirements 251
SRV records 445,937–938
unified namespaces 892,936–937
user principal names 83–84
zones 71
documentation
conventions x
Enterprise Deployment Guide described vii–viii
Enterprise Deployment Guide vs. Resource Guide vii
Resource Guide described viii–ix
Resource Kit companion CD ix
support policy ix
domain controllers
deployment considerations 209
domain naming master 879
global catalogs 878–879
infrastructure master 880
namespaces 66–68
number of 67–68
operations master 879
optimizing 877–880
PDC emulator 880
placement 396
relative identifier master 879
roles. list of 877
domain controllers *(continued)*
schema master 879
servers during upgrades 159–160
domain local groups 101–102, 833
Domain Name System *See* DNS
domain naming contexts 96
DomainPrep
preparing for mixed mode 142–144
requirements 250–251
security 859–862
domains
Active Directory design impact 75–76
administrative partitioning 874
changing after deployment 75
costs for additional 73–74
design 72–76,208–209
first in forest 67
forest root 74–75
logon names 82
mixed vs. native mode 68
namespaces 66
naming recommendations 79–82
number of 72–74, 139,874
optimizing 873–875
partitions 896
physical partitioning 875
single tree with four domains 69
structure 46–47
trees *See* trees
upgrading 139, 158–159
downtime
causes of 338
costs 335–336
DSAccess (Directory Service Access)
API overview 384

DSAccess (Directory Service Access) *(continued)*
cache 900–903
configuration 903–904
directory cache 386
settings 386–387,900
specifying Active Directory servers 384–385
DSProxy Service
client access to Active Directory 634–635
Outlook 2000 635–636
Outlook 98 124–125,635–636
specifying global catalog servers 634–635
traffic and load generated 637
dual authentication 569
dual-homed system 868
dual-key vs. single-key pairs 529–530
dynamic data 768
Dynamic RAS Connector 270

E

edb files 761
EFS (Encrypting File System) 842–843
e-mail
connectivity ***See*** connectivity
inter-company security deployment scenarios 550–552
intra-company security deployment scenarios 546–550
security 521–523
test plans 183
Encrypting File System (EFS) 842–843
encryption
algorithms 853
inbound 865–866
Key Management Service 851–852
message body 521–522
encryption ***(continued)***
RPCs (remote procedure calls) 866
enhancements ***See*** what's new
enrollments, certificate 524, 527–530, 547–548
enterprise administrators 858–859
Enterprise Deployment Guide vii–viii
entry points for viruses 320–321
environments
branch office ***See*** branch office deployment
deployment ***See*** deployment environment
hosted service ***See*** hosting other companies
low-bandwidth 395–396,447
pilot ***See*** pilot programs
security ***See*** security
testing ***See*** testing
envisioning phase
building project teams 23–24
conceptual designs 27
described 20
high-level requirements 25
overview 22
project assumptions 26
project scope 26
project structures 25
project visions 25
error messages
-1216 970–971
0xC103798A 949
0xEFAD2521 956
0xfffffb40 970–971
454 Client must be authenticated 968
ADSI 723
Conference Not Found 965
Deleted Server 966
Event ID 1018,1019, and 1022 970

error messages (continued)
Exchange 5.5 to Exchange 2000 Mailbox Move Failed 958–959
Idispatch not found 949–950
Setup failed while configuring registry entries 949
Setup was unable to bind to the Exchange Server 948–949
System cannot find the path specified 957
System error 2 has occurred 955–956
The Directory Service is busy 949
The following recipient could not be reached 967
The store could not be mounted 957–958
ESE (Extensible Storage Engine)
database consistency 784
defragmenting databases 785–786
stnrcture 760
ESEUTIL 787–789
event logs 265
event scripts 276
Event Sink Template Wizard for Visual Basic 757
Event Viewer 303–304.809–810
events
backup 801
NNTP inbound and outbound events 734–735
sinks 738
SMTP inbound and outbound events 734–735
SMTP transport 745–749
transport 15
types of 738
Web Storage System 15,738–744
examples, risk assessment 21–22
Exchange 2000 Conferencing Server See Conferencing Server
Exchange 5.5
See *also* coexistence
administration in mixed mode 297–299
administrative groups 291
administrative tools 899
architecture 700–701
backup and restore 305–309
compared Io Exchange 2000 attributes 266–269
compared to Exchange 2000 components 255–257
compared to Exchange 2000 design 50
compared to Exchange 2000 Internet Mail Connector 271–275
compared to Exchange 2000 objects 282
compared to Exchange 2000 public folders 293–294
compared to Exchange 2000 Site Connector 270
compared to Windows 2000 sites 90
components integrated into Windows 2000 57
connection agreements See connection agreements
container structures 140
custom attribute names 258
dependencies for upgrading 251–252
directory preparation 140–141
Directory Service vs. Active Directory 111
distribution lists vs. groups 120
GWART vs. link state routing 377–378
hidden objects 258
Instant Messaging deployment 440–441
Internet Mail Service vs. SMTP connector 376
inter-organization solutions 450-453
message routing basics 373–374
mixed mode message routing 379–380
moving servers 298–299
Outlook Web Access 558

Exchange 5.5 *(continued)*
recipient containers vs. organizational units 257, 285,290
replicating files to Exchange 2000 703–705
services that cannot be upgraded 171
Site Connector vs Routing Group connector 375
troubleshooting when upgrading 950–951
upgrading *See* upgrading
virus scanning support in SP3 325–327
Exchange Administration Delegation Wizard
assigning roles 292, 858
granting rights to manage servers 286
Exchange Client
MAPI directory service requests 126
Outlook 98 and Outlook 97 636
Exchange Connector for Novell GroupWise *See* Novell GroupWise
Exchange Installable File System (ExIFS) *See* IFS
Exchange Instant Messaging Service *See* Instant Messaging
Exchange Interprocess Communication Layer (ExIPC) 661–664
Exchange Key Management Service *See* Key Management Service
Exchange Lotus cc:Mail *See* Lotus cc:Mail
Exchange Lotus Notes *See* Lotus Notes
Exchange Microsoft Mail *See* Microsoft Mail
Exchange Monitoring and Status Tool 300–302
Exchange MTA (message transfer agent) *See* MTA
Exchange PROFS *See* PROFS
Exchange Resource Kit companion CD ix
Exchange SNADS *See* SNADS
ExIFS (Exchange Installable File System) *See* IFS
ExIPC (Exchange Interprocess Communication Layer) 661–664
Extensible Markup Language (XML) 14,560,728
Extensible Storage Engine (ESE)
database consistency 784
defragmenting databases 785–786
structure 760
external connectivity *See* connectivity

F

failover 761, 765
fault tolerance
Outlook Web Access 564
verifying backups 801
file header output 787
file semantics 735
file viruses 322
files
See also specific files
chat migration 275
checkpoint 763
component storage locations 684–685
database 761–762
log 763
storage group 761,765
transaction log 762, 786,788
troubleshooting last access times 957
filtering, TCP/IP 846–847
Financial Bank, Inc. 487–488
firewalls
Chat Service 436
implementing 973–975
Instant Messaging 448,928–930,939
Outlook Web Access 12,563
ports and protocols 973–975
virus protection 867–868

ForestPrep
preparing for deployment 152
preparing for mixed mode 142–143
requirements 249–250
security 859–861
forests
centrally managed 458–459
connecting to 158–159
creating domain hierarchies 79–81
deploying 89
design 208
directory replication 551–552
directory synchronization in backbone scenarios 499
first domains 67
implementing Exchange 2000 89
inter-forest scenarios 458–461
limitations 873
migrating servers 160–161
multiple trees 70
multiple-forest environments 88–89, 872
number of 47, 88
number of domains 72–74
optimizing 872–873
overview 88
peer 459–461
preparing for deployment 152–153
root domains 74–75
separate 172–174
separate organizations 171–172
single-forest environments 88, 872
synchronizing data 90
formats
HTML 728
list of 709.727
formats *(continued)*
MIME 728
overview 727
XML 728
FQDN (fully qualified domain name) 938
front-end architecture
centralized administrative model 470–471
configuring Outlook Web Access servers 565–567
connecting over the Internet 213–214
implementing availability technologies 345–347
mixed administrative model 475–476
Outlook Web Access dual authentication 569
Outlook Web Access firewalls 12, 563
Outlook Web Access overview 562,565, 697–698
Outlook Web Access pass-through authentication 569
scaling servers 886
server deployment environment 224–225
server disk configurations 367
server memory 371
server specifications 224
servers described 352
FrontPage 2000 17
full backups 771
full-text indexing
best practices 890–892
optimizing 887–892
public folders 61
testing 188
what's new I1
fully qualified domain name (FQDN) 938
functional specifications 29

G

GAL (global address list)
- hosting other companies 513
- optimizing 897
- Outlook 2000 vs. Exchange 2000 99–100
- overview 120

gap anaiysis 203–205

Gateway Address Resolution Table (GWART) 377, 393

global address list *See* GAL

global catalog servers
- accessing information 99
- attributes before and after Exchange 2000 installation 117–118
- attributes for replication 897–898
- backbone scenarios 498–499
- database sizing 112–114
- deployment environment 4849,208–209
- domain controllers 878–879
- DSProxy Service 634–635
- naming contexts 98
- number required 112–115
- object sizes 112–114
- overview 98, 111–112,408
- placement 98, 114–115,214,396
- placement in branch offices 477–478
- replication 5
- selectable field replication 115–116
- selecting attributes to replicate 117
- server sizing 114

global groups 101–102,834

globally unique identifier (GUID) 762

Group Policy
- moving mailboxes 296–297
- organizational units 87

groups
- Active Directory 100–107, 831–836
- administrative 6
- compared to distribution lists 120
- creating 106
- distribution 101, 229,293, 831–833
- domain local 101–102. 833
- global 101–102, 834
- implementation strategy 835–836
- mail-enabled naming conventions 228–229
- mail-enabled vs. distribution lists 5, 293
- mail-enabling 103
- routing *See* routing groups
- scenarios 106–107
- scope 101–103,832–833
- security 101,229,260,293,831–833
- selecting which type to use 104–106
- storage *See* storage groups
- types of 100
- universal 101–103, 834

GUID (globally unique identifier) 762

GWART (Gateway Address Resolution Table) 377, 393

H

hardware
- deployment environment 55–56
- mailbox 295
- Microsoft Hardware Compatibility List 247
- requirements 247
- server 339
- test lab 179

hash functions 764, 852

heaps 662

horizontal vs. vertical scalability 351

hosting other companies
Active Directory partitions 507–508
administration 516–517
architecture 505–506
configuration checklist 516–517
customer division overview 506–507
database recovery times for sample service 511
default public folders 515
desktop services 516
GAL (global address list) 513
mailbox storage 509–510
naming conventions 511–513
Outlook 2000 vs. Outlook Web Access 514–515
overview 505
partitioning data 507–508
public folders 513–516
services based on partitioning 509
services offered 506
services, examples of 510–511
shared services 515
summary 517
target company size 509
user alias and logon name examples 512
hot fixes 760,767
HTML (Hypertext Markup Language) 728
HTTP (Hypertext Transfer Protocol) 709, 968
Httpmon exe 564, 571
Hypertext Markup Language (HTML) 728
Hypertext Transfer Protocol (HTTP) 709,968

I

Ics.dat 963
IDataSource interface 724
IFS (Installable File System)
accessing data 677–678
active/active clustering 679
architecture 677–679
inbound and outbound process 679
MBX folder 678
public folders hierarchy 678
space allocation 679
troubleshooting 955–956
IIS (Internet Information Services)
backup prerequisites 769
client protocol support 557
Instant Messaging See Instant Messaging
overview 731–732
process 661
protocol locations 709
Web Storage System integration 729–730
what's new 16
IMAP4 (Internet Message Access Protocol version 4)
architecture 672–673
branch office deployment 481–482
compared to other clients 484–486
compared to POP3 712
deployment environment 234
overview 712
upgrading 263
upgrading clients 277
implementing
availability technologies 339–347
firewalls 973–975
inbound encryption 865–866
incremental backups 771

indexing
best practices 890–892
building indexes 889
full-text 11, 61, 188, 887–892
gather files 889–890
query processor 888
requirements 890–892
scheduled updates 890
searches 888–889
size of indexes 361
individual database backups 764
Information Roadmap xi
Information Store Integrity Checker (ISINTEG) 789–791
in-place domain upgrades 154
in-place upgrades 161–163
Installable File System See IFS
installations
See also upgrading
ADC location 138
ADC requirements 145
first time 152
Outlook Web Access 562
troubleshooting 948–952
virus scanning API DLL 325
Instant Messaging
Active Directory 934–935
addressing 932–934
backbone scenarios 492
branch offices 447
caching 932
client components 931
client name resolution 445–446
client operations 931–932
components 926–934
Instant Mescaging *(continued)*
configuring 934–935
configuring DNS 936–938
configuring server components 938 979
contact subscnptionr 932
deploying 438–441
deployment scenarios 940–944
DNS SRV records 445, 937–938
enterprise deployment 941–942
Exchange 2000-only deployment 439–440
Exchange 5 5 deployment 440–441
firewalls 448, 928–930, 939
FQDN (fully qualified domain name) 938
home servers 441–444, 933
Internet Information Server 5 0 934
ISP deployment 943–944
large company secure deployment 444
locator 928
logons 931–932
low-bandwidth connections 447
mid-sized company deployment 442
multiple e-mail domain deployment 942–943
name resolution 445–446
network address translations 448
network connectivity 446–447
node database 927–928
overview 431, 437–438, 925
property sheets 933
recommendations 448
redirection 928
relaying requests across networks 444–445
routers 441–444, 928–930
RVP protocol 927
security 448
Server Application Layer 927

Instant Messaging *(continued)*
server capacity 355
server components 926–930
server disk configurations 368–369
server limits 939–940
server locations 446447
server memory 371
server scalability 446
servers described 352
short names 938
small business deployment 441,940
standard deployment 941
traffic analysis tests 631–633
troubleshooting 964–965
unified namespaces 936–937
URLs 932–934
user addresses 934
virtual servers 939
what's new 13
Windows 2000 dependencies 934
Integrated Windows Authentication 568
intelligent routing 342–344
inter-company
SNIME deployment scenarios 550–552
SNIME design scenarios 537–540
inter-company with trusted third party
accessing directory information 554
certificate revocation list (CRL) 554–556
certificate trust list (CTL) 554
replicating directory information 554
S/MIME deployment scenarios 552–556
SNIME design scenarios 541–545
Internet
connectivity 213–214,219–220
directory services 552–556
Internet *(continued)*
mail 53–54
mail architecture 494495
Outlook Web Access 563
protocols 664
security for connections 867–869
Internet Explorer 561.734
Internet Information Services *See* IIS
Internet Mail Connectors 270–275
Internet Mail Service 376
Internet Message Access Protocol version 4 *See* IMAP4
Internet News Services 963
Internet Protocol Security (IPSec) 844–846
Internet Relay Chat (IRC) 431
InterOrg Synchronization tool 451
inter-organization solutions
ADC (Active Directory Connector) 451–453
designing Metadirectory Services 460–461
Exchange 5.5 450–453
inter-forest scenarios 458–461
Metadirectory Services architecture 456–457
Metadirectory Services management agent 457
Metadirectory Services overview 454–455
Metadirectory Services summary 461
overview 449
replication between Exchange 2000 organizations 453–454
replication vs. synchronization 450
scenarios 450
setting up Metadirectory Services 460
summary 461
intra-company
SNIME deployment scenarios 546–550
S/MIME design scenarios 532–536

IPSec (Internet Protocol Security) 844–846
IRC (Internet Relay Chat) 431
ISINTEG (Information Store Integrity Checker) 789–791

Kerberos 839–841
Key Management Service
certificate enrollments 527–530
certificate renewals 530
certificate revocations 531–532
Certificate Services 853–854
ciphers 852–853
CryptoAPI 526
disaster recovery 796–797
encryption 851–852
hash functions 852
key recovery 531
overview 525, 851
passwords 526
process flow 526
SNIME *See* S/MIME
single-key vs. dual-key pairs 529–530
upgrading 276
kms files 761

LAN deployment environment 45
Layer 3 security protection 844–845
LDAP (Lightweight Directory Access Protocol)
high and very high trust levels 545
Migration Wizard memory leak 972
overview 384,713–714
paged or non-paged results 714
S/MIME access strategies 538–540
LDAP (Lightweight Directory Access Protocol) *(continued)*
troubleshooting port numbers 950
upgrading 264
leapfrog upgrades 161–163
Lightweight Directory Access Protocol *See* LDAP
Link Crawling 332
link monitors 264
link state algorithms
advantages 914
architecture 684
backbone scenarios 491,501
compared to GWART 377–378
how link state works 391–392
overview 391
updates 392–393
what's new 8
link state tables
costs 913
maintenance 914–915
optimizing 913–915
overview 906
LitWare example overview 201
load balancing 564
Local Delivery queues 810
Local Security Authority (LSA) 830–831
log files
event logs 265
signatures 763
storage groups 761–762
tracking logs 265
transaction 762, 786–788
logging
circular 262,763
diagnostic 146, 806–808

logging *(continued)*
Protocol Logging tool 808–809
protocols for virtual servers 855–856
virus scanning 327
logical hard disk counters 814
Lotus cc:Mail
Address space tab 421
Advanced tab 421
Automatic Directory Exchange (ADE) 418
changing locations 364
configuration 420–422
connector described 417
Delivery restrictions tab 422
directory synchronization 418–419
Export container tab 422
Import container tab 421–422
name formats 418
Post office tab 420
scenario 419–420
what's new 9
Lotus Notes
Address space tab 412
addresses 228
Advanced tab 413–414
architecture for coexistence 236–238
changing locations 364
configuration 411–414
connecting with Exchange 405
connector described 409
Delivery restrictions tab 412
design summary 242–243
directory synchronization 410–411
Dirsync tab 412–413
Export container tab 413
General tab 411–412
Lotus Notes *(continued)*
Import container tab 413
messaging 410
migrating to Exchange 239–240
migration overview 234–235
services for coexistence 235
what's new 9
low-bandwidth environments 395–396,447
LSA (Local Security Authority) 830–831

M

macro viruses 322
mailbox servers
capacity 354
deployment environment 51–52,221–223
described 352
disk configurations 357–362
disk controller cache settings 359–360
disk space margin of safety 361
index size 361
memory 370
MIME content 361
network configurations 372
page file settings 360
RAID 359
specifications 222
SRS (Site Replication Service) 361–362
storage system 223
mailbox store locations 358
mailbox-enabled vs. mail-enabled objects 257
mailboxes
accessing 559
centralized administrative model 472–473
coexistence 163–164
configuration connection agreements 164–170

mailboxes *(continued)*
distributed administrative model 475
distribution lists 170–171
Group Policy 296–297
GUID 762
hardware 295
logging on 695–696
mixed administrative model 477
moving 294–297
network bandwidth 295
permissions 864–865
permissions for upgrades 170–171
restoring 782
SRS (Site Replication Service) 164–170
troubleshooting identical 970
upgrading 161–171, 259
mail-enabled vs. mailbox-enabled objects 257
mail-enabling groups 103
mail-enabling public folders 398, 681
maintenance
See also monitoring
link state tables 914–915
schedules 767
MAPI (Messaging Application Programming Interface) 727
Master Project Plan 30–31
matching objects with previous versions of Exchange 135–137
MBX folder 678
memory
chat servers 372
connector servers 370
data conferencing servers 371
front-end servers 371
full-text indexing 891
memory *(continued)*
Instant Messaging servers 371
mailbox servers 370
public folder servers 370
video conferencing servers 371
merging organizations 174
message categorizer 387,491, 669
message flow
architecture 683
component storage locations 684–685
inbound from SMTP to Web Storage System 687–689
inbound through MTA 691–692
message routing architecture 683–684
message transport 684
outbound from Web Storage System to SMTP 689–691
Windows 2000 Server 686–687
message hash 764, 852
message routing
administrative groups *See* administrative groups
architecture 683–684
backbone scenarios 490–491, 494, 503
coexistence 379–380
deployment scenarios 214–218
examples 915–921
Exchange 2000 connectors 375
Exchange 5.5 routing basics 373–374
GWART vs. link state routing 377–378
high availability 342–344
Internet Mail Service vs. SMTP connector 376
mixed mode connectivity 425
optimizing examples 915–921
optimizing overview 905–906
optimizing with group expansions 906–908

message routing *(continued)*
outside Exchange 2000 organizations 919
overview 373
point-to-point routing 916–917
rerouting mail 919–920
retries 920–921
route selection 379
routing basics 373–374
routing groups ***See*** routing groups
Site Connector vs. Routing Group connector 375
store and forward routing 917
topology 380–381
troubleshooting message flow 967–968
what's new 7–9, 374
within same server 916
X.400 connectors 376
X.400 vs. SMTP 373–374
message tracking
changing locations 365
upgrading 264–265
Message Tracking Center 302
message transfer agent ***See*** MTA
message transport
architecture 684
optimizing 904–905
testing 189
Messages Awaiting Directory Lookup queues 810
messaging
coexistence in backbone scenarios 503
Lotus Notes 410
Microsoft Mail 414
Novell GroupWise 422
topology for backbone scenarios 500–503
Messaging Application Programming Interface (MAPI) 727
Metadirectory Services
Active Directory management agent 457
architecture 456–457
designing 460–461
inter-forest scenarios 458–461
inter-orgmization solution overview 454–455
overview 427
setting up 460
summary 461
Microsoft CryptoAPI 526
Microsoft Exchange Resource Kit companion CD ix
Microsoft Exchange 2000 Conferencing Server ***See*** Conferencing Server
Microsoft Hardware Compatibility List 247
Microsoft Internet Explorer 561,734
Microsoft Internet Information Services ***See*** IIS
Microsoft Mail
Address space tab 417
Advanced tab 417
calendaring 416
changing locations 365
configuration 416–417
Connections tab 417
connector description 414
Connector MTAs tab 417
directory synchronization 414-416
Interchange tab 416
Local post office tab 417
messaging 414
what's new 9
Microsoft Management Console ***See*** MMC
Microsoft Metadirectory Services ***See*** Metadirectory Services
Microsoft Outlook 2000 ***See*** Outlook 2000
Microsoft Search Service 675

Microsoft Solutions Framework 19,205–207

Microsoft Web Storage System ***See*** Web Storage System

Microsoft Windows 2000 ***See*** Windows 2000

Microsoft Windows 2000 Server ***See*** Windows 2000 Server

migrating

See *also* upgrading

accounts to Windows 2000 154–157

connectors 425

Lotus Notes to Exchange 234–240

scenarios 154–157

servers to forests 160–161

users 34–35

Migration Wizard memory leak 972

MIME (Multipurpose Internet Mail Extensions) 728

mission statements 25

mixed branch office administrative model 475–477

mixed mode

See *also* coexistence

Active Directory Connector ***See*** ADC

administration overview 297

administrative groups 291, 297

administrative models 290

architecture 700

compared to native mode 68,253

connectivity 424–425

mailbox upgrades 163–164

message routing 379–380

moving Exchange 5.5 servers 298–299

preparing Active Directory 142–144

PROFS connectivity 426

public folders 298

replicating files from Exchange 5.5 703–705

routing groups 291, 466–467

mixed mode ***(continued)***

SNADS connectivity 426

SRS (Site Replication Service) 128–129, 164–170

MMC (Microsoft Management Console)

Event Viewer 303–304,809–810

S/MIME certificate enrollments 528–529

what's new *6*

monitoring

See *also* maintenance

ADC (Active Directory Connector) 146–147

analyzing data 820–821

analyzing results 821–823

bottlenecks 823–825

counters ***See*** counters

deployment environment 55

diagnostic logging 806–808

establishing baselines 821

Event Viewer 809–810

Exchange 2000 vs. Exchange 5.5 803

features 804–810

full-text indexing objects 892

graphs for reporting 822

Httpmon.exe 564,571

large values 822

Link Crawling 332

missing data 823

Monitoring and Status tool 804–806

Network Diagnosis tool (Netdiag) 820

Network Monitor 817–819

notifications for warning and critical states 804–806

Outlook Web Access performance 569–571

overview 803

performance 299–304

Performance Logs and Alerts 303, 571, 815–816

monitoring *(continued)*
Protocol Logging tool 808–809
Queue Viewer 810
queues 820
response times 821
spikes 822
startup events 822
status conditions 806
System Monitor *See* System Monitor
Task Manager 816
Terminal Services Client 817
thread identifiers 822
throughput 820
tools overview 811
virus scanning 332
zero values 823
Monitoring and Status tool 804–806
month-week-day backup rotation schedules 775–776
move mailbox upgrades 161–163
Move Server Wizard 174
MTA (message transfer agent)
inbound message flow 691–692
message deliveries to remote servers 393
what's new 684
X.400 vs. SMTP 373–374
multipartite viruses 322
multiple databases
architecture 675–676
implementing availability technologies 340–341
testing 187
what's new 10
multiple public folder hierarchies *See* public folders
multiple storage groups 761
Multipurpose Internet Mail Extensions (MIME) 728

N

name resolution
backbone scenarios 499
Instant Messaging 445–446
namespaces
See also DNS (Domain Name System)
Active Directory overview 63–64
architecture 66
DNS overview 70
domain controllers 66–68
domains 66
overview 64
planning DNS 65–66
separate vs. unified 937
trees 69–70
unified *See* unified namespaces
naming contexts 96–98
naming conventions
administrative groups 226
deployment considerations 209
deployment environment 225–230
Domain Name System *See* DNS
domains 79–82
hosting other companies 511–513
mail-enabled contacts 227–228
mail-enabled groups 228–229
resources 229–230
routing groups 226
servers 227
SMTP alias format 227
Windows 2000 97
native mode vs. mixed mode 68,253

NetMeeting
deployment environment 234
limitations 493
Netscape Messenger
logon and logoff tests 590–591
mail-item test analysis 614–615
mail item test details 603–608
mail-item test results for IMAP 610
mail-item test results for POP3 610
public folder test analysis 622
public folder test details 619–621
public folder test results for NNTP 622
test lab configuration 581–582
network bandwidth ***See*** bandwidth
Network Diagnosis tool (Netdiag) 820
network layer in Network Monitor 818
Network Load Balancing 566–567
Network Monitor
compared to System Monitor 817
measuring network traffic 818–819
observing resource usage 817
Network News Transfer Protocol ***See*** NNTP
network protocols ***See*** protocols
network traffic ***See*** traffic analysis
networks
branch office administrative models 468
connector server configurations 372
deployment environment 43
infrastructure 43
LAN deployment environment 45
mailbox server configurations 372
non-connected 396
public folder server configurations 372
test labs 180
networks ***(continued)***
test plan effects 181–182
traffic ***See*** traffic analysis
WAN deployment environment 44
new features
Active Directory 4–5,702
administration 5–6, 281
administration model 3
administrative groups 6
ADO support 16
architecture 700–701
ASP integration 16
CDO (Collaboration Data Objects) 14–15
CDO Workflow Objects 15
Chat Service 12
client directory access 4
client support 11
clustering 701
Conference Management 13
Conferencing Server 13–14
connectivity to other e-mail systems 9
custom solution development 14–17
Data Conferencing Provider 13
data storage 10–11
design considerations 5
Exchange 5.5 vs. Exchange 2000 4
full-text indexing 11
IIS integration 16
Instant Messaging 13
link state algorithm 8
message routing 7–9, 374
MMC (Microsoft Management Console) 6
MTA (message transfer agent) 684
multiple databases 10
OLE DB support 16

new features *(continued)*
overview 3
permissions 6
policies 6
protocol support 11
real-time collaboration 12–14
Routing Group connector 8
routing groups 7
scalability 12
server events 15
SMTP 7–9
testing 186–190
third-party conferencing support 14
Video Conferencing Provider 14
Web forms 17
Web Storage System 10–11,701
Web support 10–11
workflow 15
XML support 14
newsgroups. troubleshooting 965–966
NNTP (Network News Transfer Protocol)
architecture 673–674
inbound and outbound events 734–735
overview 711
Protocol Logging tool 808–809
transport event sinks with CDO 747–748
upgrading 263
non-connected networks 396
non-production servers, restoring data to 784
Notifications interface, Monitoring and Status tool 804–806
Novell GroupWise
Address space tab 423
changing locations 364
configuration 423–424
Novell GroupWise *(continued)*
connector described 422
Delivery restrictions tab 423
directory synchronization 422
Directory synchronization schedule tab 423
Export container tab 424
Filtering tab 424
General tab 423
Import container tab 424
messaging 422
what's new 9

O

objects
accessing Exchange 560
Active Directory 407–408
Active Directory schema 881–882
mail-enabled vs. mailbox-enabled 257
managing 139
matching with previous versions of Exchange 135–137
troubleshooting 953–954
upgrading 257–258
office space for test labs 179
OfficeVision 275
offline addresses 123,255
offline backups 775,781
offline defragmenting 785
OLE DB 16,727
online backups 773–775,781
online defragmenting 785
optimizing
Active Directory bulk changes 895
Active Directory data access 893–894
Active Directory data storage 895–898

optimizing *(continued)*
Active Directory overview 871–872
Active Directory schema 881–885
administration 898–899
DNS (Domain Name System) 892
domain controller roles 877–880
domains 873–875
forests 872–873
front-end and back-end servers 886
full-text indexing 887–892
GAL (global address list) 897
global catalog attributes 897–898
link state tables 913–915
message routing examples 915–921
message routing overview 905–906
message routing with group expansions 906–908
message transport 904–905
multiple virtual servers 886–887
organizational units 875–876
overview 871
public folders 921–924
routing group connections 908–913
routing group expansions 906–908
routing group overview 905–906
server overview 885
trees 873
trust relationships 876–877
tuning See tuning
Windows 2000 sites 880
Organizational Forms Library Public Folder 255
organizational requirements 41
organizational units
ACLs (access control lists) 86
administrative control 875
centralized administrative model 473
organizational units (continued)
changing after deployment 87
characteristics 84–85
compared to recipient containers 257, 285, 290
delegating administration 85–86
Group Policy 87
hiding objects 87
hierarchy 875
object locations 87–88
optimizing 875–876
overview 84,282–283
planning process 85–87
single domain model 876
Outlook 2000
32-bit forms 735
Active Directory data access 893
address book access 126
address book view lookup (ABVL) tests 598
address details (AD) tests 598–599
address lookup (AL) tests 597
address resolution (AR) tests 594–596
ambiguous name resolution (ANR) tests 596597
branch office deployment 479–481
built-in forms modules 735
Calendar Module 737
calendaring tests 615–618
compared to other clients 484–486
compared to Outlook Web Access 514–515
contact tests 615–618
Contacts Module 737
deployment environment 234
directory access measurement analysis 599–600
DSProxy Service 635–636
extendable forms 736
fields 735

Outlook 2000 *(continued)*
GAL (global address list) 99–100
instant collaboration 736–737
Journal Module 737
logon and logoff tests 585–588
mail-item test analysis 614–615
mail-item test details 603–608
mail-item test results 608
Outlook Forms Designer 736
overview 735
public folder test analysis 622
public folder test details 619–621
public folder test results 621
spam support 328
switching between forms and run time 735
task tests 615–618
Tasks Module 737
Terminal Services tests 623–629
test lab configuration 579
test plans 184–185
traffic parameters 594
troubleshooting client messages 960–963
upgrading clients 277
view, 735
virus scanning support 328
Visual Basic Expression Service 736
Outlook 97
Active Directory data access 894
address book view lookup (ABVL) tests 598
address details (AD) tests 598–599
address lookup (AL) tests 597
address resolution (AR) tests 594–596
ambiguous name resolution (ANR) tests 596–597
calendaring tests 615–618
contact tests 615–618
Outlook 97 *(continued)*
directory access measurement analysis 599–600
Exchange Client 636
logon and logoff tests 585–588
mail-item test analysis 614–615
mail-item test details 603–608
mail-item test results 609
MAPI directory service requests 126
public folder test analysis 622
public folder test details 619–621
public folder test results 621
task tests 615–618
test lab configuration 579
traffic parameters 594
upgrading clients 276
Outlook 98
Active Directory data access 894
address book access 126
DSProxy Service 124–125.635–636
Exchange Client 636
MAPI directory service requests 126
upgrading clients 276
Outlook Express
address name resolution (ANR) tests 600
address resolution (AR) tests 600
LDAP mode 600
logon and logoff tests 588–590
mail-item test analysis 614–615
mail-item test details 603–608
mail-item test results for IMAP 609
mail-item test results for POP 609
public folder test analysis 622
public folder test details 619–621
public folder test results for NNTP 621
test lab configuration 580–581

Outlook Web Access
accessing Exchange objects 560
accessing mailboxes 559
accessing servers 693–695
address details (AD) tests 602
address lookup (AL) tests 602
address name resolution (ANK) tests 601
address resolution (AR) tests 601
anonymous access 568
architecture 557–558,693
authentication 567–569
back-end architecture overview 562, 565
Basic authentication 568
branch office deployment 483
browsers other than Intemet Explorer 561–562
calendaring tests 615–618
clients 561–562
compared to other clients 484–486
compared to Outlook 2000 514–515
configuring back-end servers 565
configuring front-end servers 565–567
contact tests 615–618
DavEx.dll passed items 697
deployment environment 234
deployment planning 562–564
directory access measurement analysis 602
DNS entry in A (host) record 567
dual authentication 569
evolution 558
Exchange 5.5 558
fault tolerance 564
features 559–560
firewalls 12, 563
front-end architecture overview 562,565
functions 698
Outlook Web Access (continued)
Httpmon.exe 564, 571
installing 562
Integrated Windows Authentication 568
Intemet Explorer 561
Intemet security 563
Kerberos delegation of authentication 841
kiosks 698
light messaging 698
limitations 562
load balancing 564
logon and logoff tests 591–592
logons 695–696
logout 696
mailbox logons 695–696
mail-item test analysis 614–615
mail-item test details 603–608, 611–613
mail-item test results 614
migration 698
monitoring tools 570–571
name resolution 601
Network Load Balancing 566–567
opening items 696–697
option and command verbs 694–695
pass-through authentication 569
Performance Logs and Alerts 571
performance monitoring 569–571
public folder test analysis 622
public folder test details 619–621
public folder test results 622
reach browsers 561–562
roving user support 698
server components 695
SSL (Secure Sockets Layer) 563,568
System Monitor 569–571

Outlook Web Access *(continued)*
test lab configuration 582
test plans 184
upgrading 263
URL syntax 695
Web Storage System 569–570
WebDAV 560
XML 560

pass-through authentication 569
passwords, Key Management Service 526
PDC (primary domain controller) emulator 880
performance
backup 770
connectors 395
deployment environment 55
disaster recovery 780
monitoring ***See*** monitoring
monitoring tools ***See*** tools
optimizing ***See*** optimizing
troubleshooting 971–972
tuning ***See*** tuning
virus scanning impact 326–327
Performance console
Performance Logs and Alerts ***See*** Performance Logs and Alerts
System Monitor ***See*** System Monitor
Performance Logs and Alerts
Outlook Web Access 571
overview 303, 815–816
performance monitoring ***See*** monitoring
performance monitoring tools ***See*** tools
performance objects 811–813
permissions
Active Directory 469
ADC (Active Directory Connector) 863–864
backup prerequisites 768
connectors 864
mailbox upgrades 170–171
mailboxes 864–865
predefined 856–857
public folders 259–260
queues 865
recipient management 286
routing groups 864
types of 856
what's new 6
phased approach for project plans 20–22
physical hard disk counters 814
physical security 867–869
pilot programs
advantages of 194
availability 196
choosing which features to pilot 196–197
client-side features 197
costs 195
determining participants 198
documenting processes 200
infrastructure requirements 197
lab configurations 199–200
lab test lessons 200
length of 199
moving from lab to production 199–200
objectives 195–197
one-way connection agreements 195
overview 193
pre-pilot tests 33

pilot programs *(continued)*
production lessons 200
production pilots 199
role of 194–195
scalability 196
server-side features 197
user expectations 198
user pilot tests 33–34
workflow 196
pipelining 489,669–671
PKI (public key infrastructure)
between companies 538
building 534–536
high and very high trust levels 542–545
third-party relationships 541–545
planning
backup strategy 766
branch office administrative models 467–469
connectivity 403–407
disaster recovery 54–55
DNS namespace 65–66
functional specifications 29
future requirements 43
gathering information 28
global catalog servers 115
identifying resources 30
Master Project Plan 30–31
organizational units 85–87
Outlook Web Access deployment 562–564
pilot programs *See* pilot programs
planning phase described 20
project overview 27–28
project schedules 31
proof-of-concept testing 29
restore strategies 779–780
planning *(continued)*
test labs 178–180
test plans 181–186
virus protection solutions 322–324
policies
CA (certification authority) 524
recipient 287,408
system 291–292
testing 189
what's new 6
POP3 (Post Office Protocol version 3)
Active Directory 713
architecture 672–673
branch office deployment 481–482
compared to IMAP4 712
compared to other clients 484–486
deployment environment 234
overview 712
SSL encryption and decryption 713
unified namespaces 713
upgrading 263
upgrading clients 277
ports
Chat Service TCP 436
list of 973–975
troubleshooting LDAP port numbers 950
Post Office Protocol version 3 *See* POP3
pre-pilot tests 33
present branch office administrative model 468
presentation/program layer in Network Monitor 819
previous versions of Exchange *See* coexistence; Exchange 5.5
primary domain controller (PDC) emulator 880
problems *See* troubleshooting
processor configurations 353–355

production for deploying projects 35
PROFS (Professional Office System)
directory synchronization 427
mixed mode connectivity 426
upgrading 275
project plans
deploying phase 34–35
developing phase 32--34
envisioning phase 22–27
overview 19
phased approach 20–22
planning phase 27–31
summary 35
Protocol Logging tool 808–809
protocol virtual servers 262–263
protocols
ExIPC stubs 663
HTTP 709,968
IMAP4 ***See*** IMAP4
Internet 664
IPSec (Internet Protocol Security) 844–846
Kerberos 839–841
LDAP ***See*** LDAP
link propagation See *link* state *algorithms*
list of 708, 973–975
location 709
logging for virtual servers 855–856
MIME 728
NNTP ***See*** NNTP
POP3 ***See*** POP3
RVP 927
SMTP ***See*** SMTP
supporr I I
TCP/IP 846–847
upgrading 262–264
protocols ***(continued)***
WebDAV 710
proxy servers 868
Public Folder Inter-organization Replication Tool 90
public folder servers
capacity 354
deployment environment 51–52,223
described 352
disk configurations 357–362
disk controller cache settings 359–360
disk space margin of safety 361
index size 361
memory 370
MIME content 361
network configurations 372
page file settings 360
RAID 359
SRS (Site Replication Service) 361–362
storage system 223
public folder store locations 358
public folders
access control 293–294
Active Directory 398–399, 681, 922
administration 261
affinity ***See*** affinity, public folder
architecture 679
backbone scenarios 490
client support for trees 921
configuring 398
connecting to 681–683
connecting to replicas 923–924
content access on alternate servers 680
costs 682
deployment environment 59–61, 231–233
Digital Dashboard 233

public folders ***(continued)***
Exchange 5.5 vs. Exchange 2000 293–294
full-text indexing 61
hierarchies 59, 232, 397
hosting other companies 513–516
latency for trees 922
[nail-enabling 398, 681
managing 922–924
mixed mode 298
multiple store support 679
multiple trees 680–681, 921–922
optimizing 921–924
overview 396–397
permissions 259–260
referrals 231–232, 262, 924
replicas 923–924
replicating system folders 923
replication 60–61, 399, 680
support for trees 921
test plans 185
testing 188
troubleshooting replication 952
upgrading 259–262
Web Storage System 259–262
public key infrastructure (PKI)
between companies 538
building 534–536
high and very high trust levels 542–545
third-party relationships 541–545

Q

Queue Viewer 302,810
queues
advanced queuing engine 387,491,669
described 820
queues *(continued)*
ExIPC binding facility 663
Local Delivery 810
Messages Awaiting Directory Lookup 810
permissions 865

R

RAID (redundant array of independent disks)
connector servers 367
mailbox servers 359
overview 356–357
public folder servers 359
RAS (remote access service) 265
real-time collaboration
Chat Service ***See*** Chat Service
Conferencing Server ***See*** Conferencing Server
deployment environment 219–221
Instant Messaging ***See*** Instant Messaging
network configurations for servers 372
servers described 352
services described 925
what's new 12–14
recipient administrators 859
recipient containers vs. organizational units 257, 285, 290
recipient management
Active Directory Delegation of Control Wizard 285
administrative models ***See*** administrative models
administrative tools 284–285
Exchange Administration Delegation Wizard 286
organizational units 282–283, 285
overview 282
permissions 286

recipient policies 287,408
Recipient Update Service
address lists 123, 401
overview 409
polling 129
SMTP and proxy addresses 82
recovery ***See*** disaster recovery
recovery point 337–338
recovery time 337–338
redundant array of independent disks ***See*** RAID
reference databases 790
registry
editing directly 825
troubleshooting access 969
reinstalling Windows 2000 793
remote access service (RAS) 265
remote procedure calls (RPCs) 866
renewals of certificates 530
repairing databases 788
replication
Active Directory 110–111
ADC (Active Directory Connector) 137,703–705
branch office deployment 463
directory replication between two forests 551–552
Exchange 2000 inter-organization solutions 453–454
Exchange 5.5 inter-organization solutions 450–453
global catalog servers *5*
inter-organization overview 449
inter-organization replication vs. synchronization 450
latency 183
Organizational Forms Library Public Folder 255
public folders 60–61, 399,680
replication ***(continued)***
replicating files from Exchange 5.5 703–705
Schedule+ Free/Busy Public Folder 255
selectable field 115–116
selecting global catalog attributes 117
requirements
Active Directory 248–251
Active Directory schema 248–249
ADC (Active Directory Connector) 145
administrative 42
availability 334–338
backup resource 768
branch office functionality 479
branch office mobile vs. offline 479–486
business 41–43
components of Windows 2000 Server 248
dial-up access 251
disaster recovery 792
disk space 252
DNS service 251
DomainPrep 250–251
ForestPrep 249–250
full-text indexing 890–892
future planning 43
hardware 247
NNTP service 251
organizational 41
overview 246–247
permissions for upgrading 253
pilot program infrastructure 197
previous Exchange version dependencies 251–252
project plan 25
service requirements 251
SMTP service 251

requirements *(continued)*
technical 42
testing 181
time required for upgrading 252
user 42
Windows 2000 Server infrastructure 247–251
Resource Guide vii–ix
Resource Kit companion CD ix
resources
backup requirements 768
naming conventions 229–230
project planning 30
response times 821
restore ***See*** backup and restore
revocations of certificates 525,531–532
risks, project 20–22
roles for deployment teams 23–24
root certification authority 524
rotation schedules for backups 775–777
routing and transport
See also message routing
advanced queuing engine 387
connector types 388–389
distribution group expansion 388
link state information 391–393
low-bandwidth environments 395–396
message categorizer 387
message deliveries to remote servers 393
Routing Group connector 389
routing group scenario 393–395
SMTP connector 390
X.400 connector 390–391
Routing Group connector
benefits of using 911
compared to Exchange 5.5 Site Connector 375
Routing Group connector *(continued)*
compared to other connectors 908
deployment scenarios 502
overview 389
what's new 8
routing groups
backbone scenarios 500–503
boundaries 908–909
branch office deployment 464467
bridgehead servers 466,911–912
centralized administrative model 469–471
compared to sites 375
connectors, list of 906
costs for multiple 907,917–918
deployment scenarios 501–502.909–910
design considerations 908
distributed administrative model 473–474
example of 464465
grouping multiple locations 910
link state tables 906
many-to-many relationships 343
mapping sites 466
masters 254,465, 683, 914
mixed administrative model 475–476
mixed mode 291,466467
multiple 907, 917–918
multiple paths between 907
naming conventions 226
optimizing overview 905–906
optimizing with connections 908–913
optimizing with group expansions 906–908
permissions 864
physical location of users and network links 909
point-to-point routing 916–917
propagation 683

routing groups (continued)
redundant connectors 474
rerouting mail 919–920
resilience 344
retries 920–921
scenario 393–395
server management 291
SMTP connector 912–913
store and forward routing 917
testing 188–189
topology 393–394,907
traffic analysis 907
troubleshooting message flow 967–968
upgrading 254
what's new 7
RPCs (remote procedure calls) 866
running
ISINTEG 790–791
Setup in disaster recovery mode 794

S/MIME
building PKI 534–536
certificate enrollments 527–530
commercial vs. self-hosted root CA 534
deployment scenarios 546
directory access between companies 538–539
directory synchronization between companies 540
enrollments 547–548
high and very high trust levels 542–545
inter-company deployment scenarios 550–552
inter-company design scenarios 537–540
inter-company with trusted third party deployment scenarios 552–556
inter-company with trusted third party design scenarios 541–545
intra-company deployment scenarios 546–550
intra-company design scenarios 532–536
LDAP access strategies 538–540
PKI between companies 538
PKI with third party 541–545
referral access between companies 540
subdirectory access between companies 539–540
trust relationships between companies 537–540
trust relationships for internal and external clients 534–536
scalability
Chat Service 434,945
horizontal vs. vertical 351
Instant Messaging 446
pilot programs 196
what's new 12
scanning, virus
adding to monitoring services 332
administration 326
detection circumvention tests 330–331
encrypted messages 327
Exchange 5.5 SP3 325–327
installing API DLL 325
logging 327
operations 326
overview 321
performance counters 331
performance with Exchange 2000 327
performance with Exchange 5.5 326
protection without virus scanning API 327
support in Exchange 2000 Server 327
support in Outlook 328
Schedule+ Free/Busy Public Folder 255

schedules
maintenance 767
project planning 31
schema naming contexts 97
schema, Active Directory
attributes 117–119, 882
classes 881
deactivating components 883
discovery 705
extending 48, 97
modifying 882–885
network performance during replication 885
objects 881–882
optimizing 881–885
partitions 897
replication latency 884
requirements 248–249
scripting 882–883
snap-in 882
syntax rules 882
updates on existing objects 884
Web-based collaboration 730
Secure Sockets Layer ***See*** SSL
Secure/Multipurpose Internet Mail Extensions ***See*** S/MIME
security
access control ***See*** access control
Active Directory ***See*** Active Directory
Active Directory Connector ***See*** ADC
administration levels 858–859
analyzing 850
auditing 838–839
branch office administrative model 468
bridgehead servers 911
CA (certification authority) 524–525
security *(continued)*
Certificate Services 842
certificates 523–525
Chat Service 436
communication between clients and servers 865–866
Data Conferencing Provider 220
deployment scenarios ***See*** deployment scenarios
digital signatures 523
DNS (Domain Name System) 77,868–869
DomainPrep 859–862
dual-homed system 868
EFS (Encrypting File System) 842–843
e-mail 521–523
encryption ***See*** encryption
Exchange 2000 features 851
firewalls ***See*** firewalls
ForestPrep 859–861
Instant Messaging 448
Internet connections 867–869
intranet messaging 845–846
IPSec (Internet Protocol Security) 844–846
Kerberos 839–841
Key Management Service ***See*** Key Management Service
Layer 3 protection 844–845
overview 519, 827
permissions ***See*** permissions
physical 867–869
proxy servers 868
risks 827–828
S/MIME ***See*** S/MIME
Security Configuration and Analysis tool 848
Security Configuration Tool Set 847–848
Security Template tool 848–850

security *(continued)*
session-based vs. message-based 520–521
TCP/IP filtering 846–847
types of attacks 828
updates 869
virtual servers 854–856
virus protection ***See*** virus protection
Web sites 869
Windows 2000 features 829
workflow 755–756
Security Configuration and Analysis tool 848
Security Configuration Tool Set 847–848
security groups
administration 293
advantages 101
described 229,831
scope 832–833
upgrading public folders 260
Security Template tool 848–850
selectable field replication 115–116
self-hosted vs. commercial root CA 534
self-signed certification authority 524
server monitors 264
servers
See also *specific* servers
accessing Outlook Web Access 693–695
Active Directory 384–385
administration 5
architecture 333–334, 347,659–660
availability for pilot programs 196
availability overview 333
availability processes 348
availability requirements 334–338
availability summary 349
back-end ***See*** back-end architecture
servers *(continued)*
capacity 353–355
client access 58
client communication security 865–866
clusters ***See*** clustering
configurations 55–58
connecting to forests 158–159
deployment environment 55–58
design issues for upgrades 159
disk configurations 356–369
domain controllers during upgrades 159–160
events 15
front-end ***See*** front-end architecture
hardware 55–56, 339
high availability architecture 333–334
implementing availability technologies 339–347
location 210–212
management features 290–292
management overview 287
management responsibilities 287–288
memory configurations 370–372
migrating to forests 160–161
moving Exchange 5.5 298–299
names 58
naming conventions 227
network configurations 372
not upgrading 158
operating system components 57
optimizing overview 885
Outlook Web Access components 695
processor configurations 353–355
protocol virtual 262–263
proxy server security 868
recovering cluster servers 797–800
recovering member servers 793–797

servers *(continued)*
restoring data to non-production servers 784
roles 221–225
scalable and resilient architecture 347
sizing **See** sizing, server
software versions 57
third-party software 57
troubleshooting message flow 967–968
troubleshooting when adding servers 951
types of 352–353
upgrading domains 158–159
upgrading from Windows NT to Windows 2000 157–161
virtual, creating multiple 886–887
virtual, security for 854–856
virtual, troubleshooting HTTP 968
virus protection 324
Setup
disaster recovery mode 794
troubleshooting 948–952
shared cache 384
sIDHistory attribute 156–157
Simple Mail Transfer Protocol **See** SMTP
single key vs. dual-key pairs 529–530
Site Connector 269–270,375
Site Replication Service **See** SRS
sites
compared to routing and administrative groups 375
deployment environment 51
described 47–48
Exchange 5.5 vs. Windows 2000 90
links 92–93
optimizing 880
Windows 2000 topology 90–94
sites *(continued)*
Windows 2000 vs. Exchange 2000 design 93–94
sizing. server
disk configurations 356–369
global catalog 112–114
memory configurations 370–372
network configurations 372
overview 351
processor configurations 353–355
smart cards 552
SMTP (Simple Mail Transfer Protocol)
advanced command verbs 488–489
advanced queuing engine 491,669
advantages 383
alias format 227
architecture 665
bridgehead servers 912
changing locations 363
chunking 488–489,671–672
compared to Exchange 5.5 Internet Mail Service 376
compared to X.400 373–374
connector, deployment scenarios for 502
connector, optimizing 912–913
connector, purpose of 390
directories 666–667
event binding rules 748–749
event bindings 746
extensions 668
ImailMsg 668
inbound and outbound events 734–735
inbound message flow to Web Storage System 687–689
message categorizer 669
outbound message flow from Web Storage System 689–691

SMTP (Simple Mail Transfer Protocol) (continued)
overview 711
pipelining 489,669–671
process 665–666
Protocol Engine (Smtpsvc.dll) 668
protocol event sinks 747
protocol events 746
Protocol Logging tool 808–809
routing engine 669
service requirements 251
store drivers 668–669
testing connector upgrades 182
transport architecture 667, 684
transport components 746
transport event sinks with CDO 747–748
transport events 745
troubleshooting 966–967
what's new 7–8, 9
SNADS (SNA Distribution System)
directory synchronization 427
mixed mode connectivity 426
upgrading 275
soft recovery 786, 789
software
client test plans 183–185
test labs 180
third-party 57
versions 57
space allocation 734
spam 322,328
SQL
ADSI object definitions 721
Web Storage System commands 716–717
SRS (Site Replication Service)
arbitration 169–170
disabled on second Exchange 2000 server 165–166
disaster recovery 795–796
enabled on bridgehead server 166
enabled on Exchange 2000 server 165
mailbox servers 361–362
mixed mode operations 128–129, 164–170
public folder servers 361–362
troubleshooting 963
upgrading multiple sites 168–169
srs files 761
SSL (Secure Sockets Layer)
enrolling certificates 548
Outlook Web Access 563,568
POP3 encryption and decryption 713
reducing overhead 346
standardizing backup formats 767
static data 767
Status interface, Monitoring and Status tool 806
stm files 761
storage groups
compared to databases 306–307
described 761
files 761, 765
moving databases 783
testing 187
troubleshooting missing databases 970–971
troubleshooting mounting or dismounting databases 959–960
stream cipher 853
structure, company 40–41
subordinate certification authority 524

symmetric cipher 852
system drives, restoring 793
System Manager virus scanning 332
System Monitor
compared to Network Monitor 817
logical hard disk counters 814
Outlook Web Access 569–571
overview 811
performance objects 811–813
physical hard disk counters 814
when to use 303
workload balance 815
system policies 291–292
system vulnerability 336–337

tapebackups 800–801
Task Manager 816
TCP/IP filtering 846–847
technical requirements 42
Temp directory locations 365
templates, deployment planning 22
Terminal Services
branch office deployment 484–486
client overview 817
compared to other clients 484–486
Outlook 2000 tests 623–629
test lab client configuration 582
test lab configuration 579
test labs
See also test plans; testing
budgeting 178
hardware 179
networking 180
office space 179
test labs *(continued)*
overview 178
software 180
traffic analysis client characteristics 577
traffic analysis client configurations 579–582
traffic analysis configuration overview 574–575
traffic analysis Exchange 2000 configuration 578
traffic analysis LAN designs 576
traffic analysis measurement methodology 583–584
traffic analysis Netscape Messenger configuration 581–582
traffic analysis Outlook 2000 configuration 579
traffic analysis Outlook 97 configuration 579
traffic analysis Outlook Express configuration 580–581
traffic analysis Outlook Web Access configuration 582
traffic analysis server characteristics 576–577
traffic analysis Terminal Services client configuration 582
traffic analysis Terminal Services configuration 579
test plans
See also test labs; testing
administrative access 182
client software 183–185
contingencies 181–183
e-mail 183
network effects 181–182
network traffic *See* traffic analysis
Outlook 184–185
Outlook Web Access 184
overview 181
public folders 185
replication latency 183
risks 181–183

test plans *(continued)*
SMTP connector 182
test strategies 186
third-party connectors 183, 185
training 183
testing
See also test labs; test plans
ADC migration paths 191
address lists 188
administrative groups 189
antivirus solutions 329–331
backup and restore 187
backups 802
designs 32
distribution lists 190
full-text indexing 188
importance of 177–178
message transport 189
multiple databases 187
network traffic *See* traffic analysis
new features 186–190
pilot programs *See* pilot programs
policies 189
pre-pilot test 33
proof-of-concept tests 29
public folders 188
requirements 181
routing groups 188–189
scenarios 186–191
storage groups 187
test environment overview 177
upgrades 190–191
user pilot tests 33–34
Web folders 190
third-party
antivirus solutions 329–332
conferencing support 14
connector test plans 183, 185
high and very high trust levels 542–545
PKI (public key infrastructure) 541–545
S/MIME deployment scenarios 552–556
S/MIME design scenarios 541–545
software 57
upgrading software 276
threads 836
three-tier virus protection 322–323
throughput 820
Time To Live setting (CacheTTL) 901
tools
See also wizards
Active Directory Migration 155–157
Active Directory Sizer 112
administrative 284–285
chat migration 275
ESEUTIL 787–789
Event Viewer 303–304,809–810
Exchange Monitoring and Status 300–302
Httpmon.exe 564, 571
InterOrg Synchronization 451
ISINTEG 789–791
Message Tracking Center 302
Monitoring and Status 804–806
monitoring tools overview 811
Network Diagnosis (Netdiag) 820
Network Monitor 817–819
Performance Logs and Alerts 303, 571,815–816
Protocol Logging 808–809
Public Folder Inter-organization Replication 90

tools *(continued)*
Queue Viewer 302,810
reactive vs. proactive 300
Resource Kit companion CD ix
Security Configuration and Analysis 848
Security Configuration Tool Set 847–848
Security Template 848–850
System Monitor See System Monitor
Task Manager 816
Terminal Services Client 817
troubleshooting Windows 2000 Server 969
Web Storage System Schema Designer 757
Web Storage System Viewer 756
Windows 2000 284,302–304
top-level hierarchies *See* public folders
topology
directory for backbone scenarios 496–499
Lotus Notes and Exchange coexistence 237
message routing 380–381
messaging for backbone scenarios 500–503
routing groups 393–394,907
Windows 2000 site 90–94
Towers of Hanoi backup rotation scheme 776–777
TP4 (Transport Class 4) 265
tracking logs 265
traffic analysis
address book view lookup (ABVL) tests 593
address details (AD) tests 593
address lookup (AL) tests 593
address resolution (AR) tests 593
ambiguous name resolution (ANR) tests 593
calendaring tests 615–618
capture filter 583
client characteristics for test labs 577
client configurations for test labs 579–582
traffic analysis *(continued)*
client test conclusions 633
contact tests 615–618
data capture 583–584
directory access tests 592–593
DSProxy chent access to Active Directory 634–635
DSProxy traffic and load 637
Exchange 2000 test lab configuration 578
Instant Messaging tests 631–633
LAN test lab designs 576
logon and logoff tests described 584–585
mail-item test analysis 614–615
mail-item test details 603–608
measurement methodology 583–584
measuring network traffic 818–819
Netscape Messenger IMAP mail-item test results 610
Netscape Messenger logon and logoff tests 590–591
Netscape Messenger NNTP public folder test results 622
Netscape Messenger POP mail-item test results 610
Netscape Messenger test lab configuration 581–582
Outlook 2000 ABVL tests 598
Outlook 2000 AD tests 598–599
Outlook 2000 AL tests 597
Outlook 2000 ANR tests 596–597
Outlook 2000 AR tests 594–596
Outlook 2000 directory access measurement analysis 599–600
Outlook 2000 DSProxy Service 635–636
Outlook 2000 logon and logoff tests 585–588
Outlook 2000 mail-item test results 608
Outlook 2000 public folder test results 621

traffic analysis *(continued)*
Outlook 2000 test lab configuration 579
Outlook 2000 traffic parameters 594
Outlook 2000 with Terminal Services tests 623–629
Outlook 97 ABVL tests 598
Outlook 97 AD tests 598–599
Outlook 97 AL tests 597
Outlook 97 ANR tests 596–597
Outlook 97 AR tests 594–596
Outlook 97 directory access measurement analysis 599–600
Outlook 97 Exchange Client 636
Outlook 97 logon and logoff tests 585–588
Outlook 97 mail-item test results 609
Outlook 97 public folder test results 621
Outlook 97 test lab configuration 579
Outlook 97 traffic parameters 594
Outlook 98 DSProxy Service 635–636
Outlook 98 Exchange Client 636
Outlook Express ANR tests 600
Outlook Express AR tests 600
Outlook Express IMAP mail-item test results 609
Outlook Express LDAP mode 600
Outlook Express logon and logoff tests 588–590
Outlook Express NNTP public folder test results 621
Outlook Express POP mail-item test results 609
Outlook Express test lab configuration 580–581
Outlook Web Access AD tests 602
Outlook Web Access AL tests 602
Outlook Web Access ANR tests 601
Outlook Web Access AR tests 601
Outlook Web Access directory access measurement analysis 602
traffic analysis *(continued)*
Outlook Web Access logon and logoff tests 591–592
Outlook Web Access mail-item test results 614
Outlook Web Access name resolution 601
Outlook Web Access public folder test results 622
Outlook Web Access test lab configuration 582
overview 573–574
public folder test analysis 622
public folder test details 619–621
routing groups 907
server characteristics for test labs 576–577
task tests 615–618
Terminal Services test lab client configuration 582
Terminal Services test lab configuration 579
test lab configuration overview 574–575
Web Storage System tests 630–631
transaction log files
database consistency 786
database technology 762
headers 788
Transport Class 4 (TP4) 265
transport events 15
transport layer in Network Monitor 818–819
transport stacks 265–275
trees
client support for public folder 921
latency for public folder 922
multiple public folder 680–681, 921–922
multiple within forest 70
namespaces 69–70
optimizing 873
single with four domains 69
support for public folder 921

trojan horse viruses 322
troubleshooting
See also error messages; monitoring
Active Directory 953–954
Active Directory Users and Computers 953
ADC (Active Directory Connector) 954–955
adding new servers 951
backup and restore 970–971
clients displaying incorrect public folders 959
connectivity 960–968
HTTP virtual servers 968
IFS (Installable File System) 955–956
installations 948–952
Instant Messaging 964–965
Internet News Service Ics.dat file 963
last access time on files 957
LDAP Migration Wizard memory leak 972
LDAP port numbers 950
mailboxes that are identical 970
message flow between routing groups and servers 967–968
newsgroups 965–966
objects 953–954
Outlook 2000 client messages 960–963
overview 947
performance 971–972
public folder replication 952
registry access 969
Setup 948–952
SMTP 966–967
SRS (Site Replication Service) 963
storage groups mounting or dismounting databases 959–960
storage groups with missing databases 970–971
upgrading from Exchange 5.5 950–951
troubleshooting *(continued)*
URL unavailable during OnCreate event 952
viruses *See* virus protection
Web Storage System 955–960
Windows 2000 Server tools 969
trust relationships
authentication with shortcut 96
between companies 537–540
certificate revocation list (CKL) 545
internal and external clients 534–536
LDAP 545
limiting 95, 877
non-transitive trust 94–95, 876
optimizing 876–877
overview 94
shortcut 96
transitive trust 94–95,877
trusted third-party relationships 541–545,552–556
tuning
address lists 400–401
directory access 384–387
DSAccess (Directory Service Access) 900–904
overview 383,900
public folders 396–399
relocating database, log, and stm files 904
routing and transport 387–396

U

unified namespaces
compared to separate namespaces 937
configuring 82–83
DNS (Domain Name System) 892
front-end and back-end architecture 345
Instant Messaging DNS 936–937

unified namespaces *(continued)*
POP3 server administration 713
universal groups 101–103, 834
upgrading
See also deployment strategies
Active Directory Connector *See* ADC
attributes 266–269
Chat Service 275
clients 276–277
coexistence with previous mailbox versions 163–164
component comparison 255–257
configuration connection agreements for mailboxes 164–170
connection agreements *See* connection agreements
connectors 265–275
design issues for servers 159
disk space requirements 252
distribution lists 170–171, 260
domain controllers 159–160
domains 139, 158–159
Dynamic RAS Connector 270
event scripts 276
Exchange 5.5 vs. Exchange 2000 components 255–257
hardware requirements 247
IMAP4 263
IMAP4 clients 277
in-place domains 154
Internet Mail Connector 270–275
Key Management Service 276
LDAP 264
link monitors 264
mailboxes 161–171,259
upgrading *(continued)*
matching objects with previous versions of Exchange 135–137
message tracking 264–265
method advantages and disadvantages 162–163
migrating sewers to forests 160–161
migration scenarios 154–157
mixed mode operations with SRS 128–129,164–170
mixed vs. native mode 253
NNTP 263
objects 257–258
OfficeVision 275
offline address books 255
operating system dependencies 246
Organizational Forms Library Public Folder 255
Outlook 2000 clients 277
Outlook 98 and Outlook 97 clients 276
Outlook Web Access 263
overview 245
permissions for mailboxes 170–171
permissions required 253
POP3 263
POP3 clients 277
previous Exchange version dependencies 251–252
PROFS 275
protocols 262–264
public folders 259–262
RAS (remote access service) 265
requirements overview 246–247
routing groups 254
Schedule+ Free/Busy Public Folder 255
server monitors 264
server roles 253

upgrading *(continued)*
servers not to upgrade 158
services that cannot be upgraded 171
Site Connector 269–270,375
SMTP connector 182
SNADS 275
steps 152–153
tests 190–191
third-party software 276
time required 252
TP4 (Transport Class 4) 265
transport stacks 265–275
troubleshooting 950–951
Web Storage System 259–262
Windows 2000 Server infrastructure 247–251
Windows NT servers 157–161
X.400 connectors 266–269
URL unavailable during OnCreate event 952
user connection agreements 131–132
user pilot tests 33–34
user principal names 83–84
users
administration 5
migrating 34–35
requirements 42
utilities *See* tools

vandals 322
verifying backups 801
version compatibility for address lists 401
vertical vs. horizontal scalability 351
Video Conferencing Provider 14
video conferencing servers
capacity 355
described 352
disk configurations 368
memory 371
virtual servers
creating multiple 886–887
security 854–856
troubleshooting HTTP 968
virus protection
adding scanning to monitoring services 332
alerts 332
antivirus checklists 329
antivirus vendors 329–332
attachments with viruses 330
backbone infrastructure 323–324
connector server antivirus solutions 367
costs 319
definitions 322
desktops 324
detection circuinvention tests 330–331
encrypted messages when scanning 327
entry points 320–321
firewalls 867–868
installing virus scanning API DLL 325
Link Crawling 332
local servers 324
logging during scanning 327
overview 319–321
performance counters for scanning 331
performance during scanning with Exchange2000 327
performance during scanning with Exchange 5.5 326

virus protection *(continued)*
scanning administration 326
scanning concepts 321
scanning operations 326
scanning support in Exchange 2000 Server 327
scanning support in Exchange 5.5 SP3 325–327
scanning support in Outlook 328
solutions planning 322–324
testing antivirus solutions 329–331
three-tier 322–323
user practices 867
without virus scanning API 327
vision statements 25
vulnerability. system 336–337

W

WAN deployment environment 44
Web Client S/MIME certificate enrollments 528–529
Web Distributed Authoring and Versioning ***See*** WebDAV
Web folders 190
Web fonns 17
Web sites
Active Directory client extensions 441
Active Directory Sizer tool 112
circular logging 262
cryptographic service providers 526
Digital Dashboard 233
DNS zones 71
E-mail Attachment Security Update 328
Exchange 5.5 backup and restore 305
Exchange 5.5 SP3 252
Exchange Up-to-Date vii
ForestPrep information 250
hosting other companies 516
Web sites *(continued)*
hot fixes 760
Httpmon.exe 564
InterOrg Synchronization tool 451
Metadirectory Services 90, 461
Microsoft Hardware Compatibility List 247
Microsoft Office macro viruses 328
Outlook 2000 vs. Outlook Web Access 514
Outlook 98 Archive Patch 592
routing masters 254
security 869
Web Storage System
accessing items 743–744
architecture 661, 674
asynchronous events 742
backbone scenarios 489–490
content storage 676–677
data center installations 316
evenls overview 738–739
events, what's new 15
HTTP functions 709
IIS integration 729–730
inbound message flow from SMTP 687–689
large department servers 313
mailbox store locations 358
Microsoft Search Service 675
Move Mailbox function 259
multiple database support 675–676
native content storage 676
navigating with ADO 714
on-demand content conversion 677
outbound message flow to SMTP 689–691
Outlook Web Access 569–570
overview 729
public folder store locations 358

Web Storage System *(continued)*
public folders 259–262
querying with ADO 716–717
recovering databases 795
registering events 743
restoring 783
rich-text store 677
sample configurations 357–358
schema 729
simple design 311
SQL commands 716–717
storage groups See storage groups
storage locations 675
streaming capability 12
synchronous events 739–742
traffic analysis tests 630–631
troubleshooting 955–960
upgrading 259–262
virus scanning 325–326
WebDAV functions 710
what's new 10–11,701
Web Storage System Application Deployment Wizard 757
Web Storage System Schema Designer 757
Web Storage System Viewer 756
Web support 10–11
Web-based collaboration
Active Directory schema 730
API components 714–727
applications, list of 731
ASP 732–733
formats 727–728
IIS 731–732
Internet Explorer 734
network protocols 709–714
Web-based collaboration *(continued)*
NNTP inbound dnd outbound events 734–735
Outlook 2000 735–737
overview 708–709
SMTP inbound and outbound events 734–735
Web Storage System 729–730
Windows 2000 Server 731
WebDAV (Web Distributed Authoring and Versioning)
accessing Inbox items 699
architecture 698–699
features 710
functions used with Web Storage System 710
Outlook Web Access 560
what's new
Active Directory 4–5, 702
administration 5–6, 281
administration model 3–9
administrative groups 6
ADO support 16
architecture 700–701
ASP integration 16
CDO (Collaboration Data Objects) 14–15
CDO Workflow Objects 15
Chat Service 12
client directory access 4
client support 11
clustering 701
Conference Management 13
Conferencing Server 13–14
connectivity to other e-mail systems 9
custom solution development 14–17
Data Conferencing Provider 13
data storage 10–11
design considerations 5
Exchange 5.5 vs. Exchange 2000 4

what's new (continued)
full-text indexing 11
IIS integration 16
Instant Messaging 13
link state algorithm 8
message routing 7–9, 374
MMC (Microsoft Management Console) 6
MTA (message transfer agent) 684
multiple databases 10
OLE DB support 16
overview 3
permissions 6
policies 6
protocol support 11
real-time collaboration 12–14
Routing Group connector 8
routing groups 7
scalability 12
server events 15
SMTP 7–9
testing new features 186–190
third-party conferencing support 14
Video Conferencing Provider 14
Web forms 17
Web Storage System 10–11, 701
Web support 10–11
workflow 15
XML support 14
Windows 2000
ADC (Active Directory Connector) 128
branch office issues 463–467
centralized administrative model 472–473
compared to Exchange 5.5 sites 90
dependencies 706
distributed administrative model 475
Windows 2000 ***(continued)***
domain structure 46–47
Exchange 5.5 components 57
Instant Messaging 934
mail-enabled groups 5
migrating accounts 154–157
mixed administrative model 477
naming conventions 97
optimizing sites 880
reinstalling 793
restoring system state 794
security features 829
site topology 90–94
sites 47–48
tools 284, 302–304
Windows 2000 Backup
IIS backup prerequisites 769
restoring Key Management Service databases 797
restoring single databases 783
restoring Web Storage System databases 795
restoring Windows 2000 system state 794
verifying backups 801
Windows 2000 Professional 285
Windows 2000 Server
Data Conferencing Provider 945
described 731
infrastructure 247–251
message flow 686–687
troubleshooting tools 969
upgrading from Windows NT 157–161
Windows 2000 Terminal Services ***See*** Terminal Services
Windows NT 4.0
domain structure 46–47
upgrading servers to Windows 2000 157–161

wizards
Active Directory Cleanup 154–155
Active Directory Delegation of Control 285
Event Sink Template 757
Exchange Administration Delegation 286, 292, 858
Migration 972
Move Server 174
Web Storage System Application Deployment 757
workflow
CDO event sinks 751–752
components 750
creating processes 750–751
described 750
engine 752–754
pilot programs 196
Restricted Mode 750
role memberships 750
security 755–756
what's new 15
Workflow System Account 754–755
workload balance 815
World Wide Web Distributed Authoring and Versioning *See* WebDAV
worms 322

X.25 connector 502
X.400 connector
Administrative Management Domain Name 269
changing locations 364
compared to SMTP 373–374
deployment scenarios 502
message routing 376
routing and transport 390–391
X 400 connector *(continued)*
upgrading 266–269
what's new 9
X.25 502
XML (Extensible Markup Language) 14,560,728

Z

zones 71