




Clifton Leonard
Brian Svidergol
Byron Wright
Vladimir Meloski

Mastering Microsoft® Exchange Server 2016



Mastering Microsoft® Exchange Server 2016



Mastering Microsoft® Exchange Server 2016

Clifton Leonard

Brian Svidergol

Byron Wright

Vladimir Meloski

 **SYBEX®**
A Wiley Brand

Senior Acquisitions Editor: Kenyon Brown
Development Editor: Kelly Talbot
Technical Editor: Joseph Nguyen
Production Editor: Athiyappan Lalith Kumar
Copy Editor: Kathy Grider-Carlyle
Editorial Manager: Mary Beth Wakefield
Production Manager: Kathleen Wisor
Executive Editor: Jim Minatel
Proofreader: Nancy Bell
Indexer: Nancy Guenther
Project Coordinator, Cover: Brent Savage
Cover Designer: Wiley
Cover Image: ©i3d/Shutterstock

Copyright © 2016 by John Wiley & Sons, Inc., Indianapolis, Indiana

Published simultaneously in Canada

ISBN: 978-1-119-23205-6
ISBN: 978-1-119-23208-7 (ebk.)
ISBN: 978-1-119-23207-0 (ebk.)

Manufactured in the United States of America

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 646-8600. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permissions>.

Limit of Liability/Disclaimer of Warranty: The publisher and the author make no representations or warranties with respect to the accuracy or completeness of the contents of this work and specifically disclaim all warranties, including without limitation warranties of fitness for a particular purpose. No warranty may be created or extended by sales or promotional materials. The advice and strategies contained herein may not be suitable for every situation. This work is sold with the understanding that the publisher is not engaged in rendering legal, accounting, or other professional services. If professional assistance is required, the services of a competent professional person should be sought. Neither the publisher nor the author shall be liable for damages arising herefrom. The fact that an organization or Web site is referred to in this work as a citation and/or a potential source of further information does not mean that the author or the publisher endorses the information the organization or Web site may provide or recommendations it may make. Further, readers should be aware that Internet Web sites listed in this work may have changed or disappeared between when this work was written and when it is read.

For general information on our other products and services or to obtain technical support, please contact our Customer Care Department within the U.S. at (877) 762-2974, outside the U.S. at (317) 572-3993 or fax (317) 572-4002.

Wiley publishes in a variety of print and electronic formats and by print-on-demand. Some material included with standard print versions of this book may not be included in e-books or in print-on-demand. If this book refers to media such as a CD or DVD that is not included in the version you purchased, you may download this material at <http://booksupport.wiley.com>. For more information about Wiley products, visit www.wiley.com.

Library of Congress Control Number: 2016946244

TRADEMARKS: Wiley, the Wiley logo, and the Sybex logo are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates, in the United States and other countries, and may not be used without written permission. Microsoft is a registered trademark of Microsoft Corporation. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc. is not associated with any product or vendor mentioned in this book.

10987654321

This book is dedicated to my loving, gorgeous wife, Marie, and to my incredible inspirations Pierce, Treyden, Gabrielle, Cheyenne, Taylor, Zoe, and Talon. Thank you for enduring all my late nights and continuously encouraging me through this journey. I love you all!

—Clifton Leonard

I'd like to thank my wife, Lindsay; my son, Jack; and my daughter, Leah, for the unending support and David Elfassy for reaching out to me to get involved with this project—thank you! Finally, I'd like to thank the original Exchange “super team”— Larry, Mike, Carl, George, Dennis, and the Chicago crew—you guys helped me elevate my game.

—Brian Svidergol

I dedicate this book to my parents who unwittingly put me on the path to working with technology by indulging me in my youth. Who knew buying a Commodore VIC-20 would get it all started? I am thankful for that and your support in many other ways over the years.

—Byron Wright

To my loving family who always supports me.

—Vladimir Meloski

Acknowledgments

Thank you once again, Microsoft, for a great release of Exchange Server. This is now the eighth major release of the well-known premier messaging system. In this release, we can see the effort and ingenuity come together in solving customer problems to create a truly superior product. Congratulations!

As the team that is working on this book completes the final steps required to send it to the printer, I continue to bring some real-world expertise into the content. I have deployed several Exchange Server 2016 infrastructures to date, but this product is so vast and so broad that I continue to find design options, best practices, and architecture recommendations on a daily basis. I'm pretty sure that I will be updating the content up to the last minute!

When I was approached to take on this book, several months before Exchange Server 2016 was about to release to manufacturing, my reaction was, "What about David and Jim?" David Elfassy authored the previous edition and has been an invaluable contributor to the Microsoft, and more specifically Exchange Server, community. Prior to David, Jim McBee authored three previous editions of this book and has been the pillar of the *Mastering Exchange Server* series. I consider it to be a true honor to take over for David Elfassy and Jim McBee as the lead author for this book, and I hope that this edition has adequately followed through on their traditions.

Throughout the book, we have tried to keep the tone and language similar to what was used in the previous editions of this book, so if you are familiar with both of these men's writing style, you should find comfort in these pages. In addition, we have removed some of the introductory technical information from previous editions, to reflect the depth of initial experience of the readers.

Taking on the responsibility of a 816-plus-page manual is no simple task and not one that can be undertaken by only one person. Along the way, I have invited several contributors to this effort. Their knowledge and expertise have added incredible value to this book. Having written anywhere from several paragraphs to complete chapters, Brian Svidergol, Byron Wright, and Vladimir Meloski are Exchange Server gurus who have provided key content for this book. These men are well respected within the Exchange Server community and are authors of Microsoft Official Curriculum, including Exchange Server 2016. They have been great contributions to this effort. Thank you!

There is also a man who has kept us all honest and has been the gatekeeper for technical accuracy in this book, and he has helped revise a couple of chapters more substantially. Joseph Nguyen agreed to take on the responsibility of technical reviewer for this book and has done a formidable job. I consider it an honor to have worked with him! Joseph, thank you!

The great folks at Wiley have been patient beyond belief when it comes to deadlines, content, and outline changes as well as our ever-changing list of contributors. They include acquisitions editor Ken Brown, developmental editor Kelly Talbot, and production editor Athiyappan Lalith Kumar.

And a special acknowledgment to those in my daily life, my father, DC Leonard; my mother, Lynette Leonard; my sister, Jaena Poppe; and my brothers, Jerry, Adam, and Jeff: thank you for always being supportive of all my endeavors.

—Clifton Leonard

About the Authors

Clifton Leonard, MCSE: Exchange Server, has more than 25 years' experience in the IT industry as an engineer, architect, consultant, trainer, and author. Clifton has extensive experience consulting on Active Directory, Exchange Server, Lync and Skype for Business Server, Identity Management, Office 365, and Azure cloud solutions. His clients include large energy corporations, K-12 schools, universities, technology manufacturers, financial institutions, the United States Air Force, and the Department of Defense. While Clifton cut his teeth on Microsoft Mail on Novell Netware and Exchange Server 5.0 on DEC Alpha, he has worked with every version of Exchange Server since then. He has also contributed as a subject matter expert to multiple Microsoft courses including Windows Desktop, Windows Server, Exchange Server, SharePoint Server, HyperV, Identity Management, Office 365, and Azure. Helping organizations migrate to the latest versions of Microsoft Exchange Server has always been a key focus of Clifton's consulting commitments.

Brian Svidergol builds Microsoft infrastructure and cloud solutions with Windows, Microsoft Exchange, Active Directory, Office 365, and related technologies. He holds the Microsoft Certified Trainer (MCT), Microsoft Certified Solutions Expert (MCSE) – Server Infrastructure, and several other Microsoft and industry certifications. Brian has authored books on Active Directory, Windows Server, Exchange Server, and related infrastructure technologies. He served as an MCT Ambassador at TechEd North America 2013 and at Microsoft Ignite 2015. Brian works as a subject matter expert (SME) on many Microsoft Official Curriculum courses, edX courses, and Microsoft certification exams. He has authored a variety of training content, blog posts, and practice test questions and has been a technical reviewer for a large number of books.

Byron Wright is the owner of BTW Technology Solutions where he provides, designs, and implements solutions using Exchange Server and Office 365. He has been a consultant, author, and instructor for 20 years, specializing in Exchange Server, Windows Server, Office 365, network design, network security, and related technologies. Byron has been a Microsoft MVP for Exchange Server since 2012.

Vladimir Meloski is a Microsoft Most Valuable Professional on Office Server and Services, Microsoft Certified Trainer, and consultant, providing unified communications and infrastructure solutions based on Microsoft Exchange Server, Skype for Business, Office 365, and Windows Server. With a bachelor's degree in computer sciences, Vladimir has devoted more than 20 years of professional experience to information technology. Vladimir has been involved in Microsoft conferences in Europe and in the United States as a speaker, moderator, proctor for hands-on labs, and technical expert. He also has been involved as an author and technical reviewer for Microsoft official courses, including Exchange Server 2016, 2013, 2010, and 2007; Office 365; and Windows Server 2012. As a skilled IT professional and trainer, Vladimir shares his best practices, real-world experiences, and knowledge with his students and colleagues and is devoted to IT community development by collaborating with IT Pro and developer user groups worldwide.

About the Technical Editor

Joseph Nguyen is a senior consultant for Microsoft. He has 20 years of experience as a system administrator, messaging engineer, IT analyst, systems engineer, consultant, and trainer providing messaging, communications, and collaboration expertise for a wide range of corporations and institutions. Joseph coauthored *Exchange Server 2010 Administration: Real World Skills for MCITP Certification and Beyond* and *MCITP Self-Paced Training Kit (Exam 70-238): Deploying Messaging Solutions with Microsoft Exchange Server 2007*.

Contents at a Glance

| | |
|---------------------------|--------------|
| <i>Introduction</i> | <i>xxvii</i> |
|---------------------------|--------------|

Part 1 • Exchange Fundamentals1

| | |
|--|-----|
| Chapter 1 • Putting Exchange Server 2016 in Context | 3 |
| Chapter 2 • Introducing the Changes in Exchange Server 2016 | 25 |
| Chapter 3 • Understanding Availability, Recovery, and Compliance | 41 |
| Chapter 4 • Virtualizing Exchange Server 2016. | 71 |
| Chapter 5 • Introduction to PowerShell and the Exchange Management Shell. | 91 |
| Chapter 6 • Understanding the Exchange Autodiscover Process. | 145 |

Part 2 • Getting Exchange Server Running 169

| | |
|---|-----|
| Chapter 7 • Exchange Server 2016 Quick Start Guide | 171 |
| Chapter 8 • Understanding Server Roles and Configurations | 195 |
| Chapter 9 • Exchange Server 2016 Requirements | 209 |
| Chapter 10 • Installing Exchange Server 2016 | 227 |
| Chapter 11 • Upgrades and Migrations to Exchange Server 2016 or Office 365 | 249 |

Part 3 • Recipient Administration 273

| | |
|---|-----|
| Chapter 12 • Management Permissions and Role-Based Access Control | 275 |
| Chapter 13 • Basics of Recipient Management. | 315 |
| Chapter 14 • Managing Mailboxes and Mailbox Content | 331 |
| Chapter 15 • Managing Mail-Enabled Groups, Mail Users, and Mail Contacts | 389 |

| | |
|---|------------|
| Chapter 16 • Managing Resource Mailboxes | 413 |
| Chapter 17 • Managing Modern Public Folders | 435 |
| Chapter 18 • Managing Archiving and Compliance | 463 |
| Part 4 • Server Administration | 493 |
| Chapter 19 • Creating and Managing Mailbox Databases. | 495 |
| Chapter 20 • Creating and Managing Database Availability Groups | 519 |
| Chapter 21 • Understanding the Client Access Services | 559 |
| Chapter 22 • Managing Connectivity with Transport Services | 597 |
| Chapter 23 • Managing Transport, Data Loss Prevention, and Journaling Rules. | 635 |
| Part 5 • Troubleshooting and Operating | 665 |
| Chapter 24 • Troubleshooting Exchange Server 2016 | 667 |
| Chapter 25 • Backing Up and Restoring Exchange Server | 691 |
| Appendix • The Bottom Line. | 719 |
| <i>Index</i> | 753 |

Contents

| | |
|-------------------------------|--------------|
| <i>Introduction</i> | <i>xxvii</i> |
|-------------------------------|--------------|

| | |
|---|----------|
| Part 1 • Exchange Fundamentals | 1 |
|---|----------|

| | |
|--|----------|
| Chapter 1 • Putting Exchange Server 2016 in Context | 3 |
|--|----------|

| | |
|--|----|
| Email's Importance | 3 |
| How Messaging Servers Work | 4 |
| What Is Exchange Server? | 5 |
| About Messaging Services | 6 |
| Many Modes of Access | 6 |
| How Messaging Services Are Used | 7 |
| The Universal Inbox | 8 |
| Architecture and Core Functionality Overview | 9 |
| Controlling Mailbox Growth | 12 |
| Personal Folders or PST Files | 13 |
| Email Archiving | 13 |
| Public Folders | 15 |
| Things Every Email Administrator Should Know | 16 |
| A Day in the Life of the Email Administrator | 16 |
| Finding Answers | 20 |
| Tools You Should Know | 23 |
| The Bottom Line | 24 |

| | |
|--|-----------|
| Chapter 2 • Introducing the Changes in Exchange Server 2016 | 25 |
|--|-----------|

| | |
|---|----|
| Getting to Know Exchange Server 2016 | 25 |
| Exchange Server Architecture | 27 |
| Windows Server 2012 R2 and Exchange Server 2016 | 27 |
| Server Roles | 28 |
| Client Connectivity | 32 |
| Hybrid Improvements | 32 |
| OneDrive for Business Integration | 33 |
| Performance | 33 |
| Improved Policy and Compliance Features | 33 |
| Message Transport Rules | 36 |
| New and Improved Outlook on the Web | 36 |
| Overview of Changes Since Exchange Server 2013 | 37 |
| Now, Where Did That Go? | 37 |
| Features No Longer Included | 37 |

| | |
|----------------------------------|----|
| Clearing Up Some Confusion | 38 |
| The Bottom Line..... | 39 |

Chapter 3 • Understanding Availability, Recovery, and Compliance41

| | |
|---|----|
| Changing from a Technology to a Business Viewpoint..... | 41 |
| What's in a Name? | 42 |
| Backup and Recovery..... | 42 |
| Disaster Recovery | 45 |
| Location, Location, Location..... | 47 |
| Management Frameworks..... | 49 |
| A Closer Look at Availability..... | 53 |
| Storage Availability | 57 |
| An Overview of Exchange Storage | 57 |
| Direct Attached Storage..... | 59 |
| Storage Area Networks | 63 |
| Compliance and Governance..... | 65 |
| The Bottom Line..... | 68 |

Chapter 4 • Virtualizing Exchange Server 2016.....71

| | |
|---|----|
| Virtualization Overview..... | 71 |
| Terminology..... | 72 |
| Understanding Virtualized Exchange..... | 73 |
| Understanding Your Exchange Environment | 75 |
| Effects of Virtualization | 76 |
| Environmental Impact..... | 76 |
| Space Impact | 76 |
| Complexity Impact..... | 77 |
| Additional Considerations | 77 |
| Virtualization Requirements | 78 |
| Hardware Requirements..... | 78 |
| Software Requirements..... | 81 |
| Operations..... | 82 |
| Deciding When to Virtualize..... | 82 |
| Deciding What to Virtualize | 83 |
| Exchange Roles | 85 |
| Testing..... | 85 |
| Possible Virtualization Scenarios | 86 |
| Small Office/Remote or Branch Office | 86 |
| Site Resilience | 87 |
| Mobile Access | 88 |
| The Bottom Line..... | 89 |

Chapter 5 • Introduction to PowerShell and the Exchange Management Shell91

| | |
|--|----|
| Why Use PowerShell? | 91 |
| Understanding the Command Syntax | 92 |
| Verbs and Nouns..... | 93 |

| | |
|--|-----|
| Help..... | 95 |
| The <i>-Identity</i> Parameter | 95 |
| Cmdlet Parameters..... | 96 |
| Tab Completion..... | 98 |
| Alias | 98 |
| Object-Oriented Use of PowerShell..... | 99 |
| Filtering Output | 100 |
| Formatting Output..... | 117 |
| Directing Output to Other Cmdlets..... | 119 |
| PowerShell v3, v4, and v5 | 120 |
| Remote PowerShell..... | 120 |
| Tips and Tricks | 121 |
| Managing Output..... | 121 |
| Running Scripts | 125 |
| Running Scheduled PowerShell Scripts | 127 |
| Debugging and Troubleshooting from PowerShell..... | 127 |
| Getting Help | 128 |
| Getting Tips..... | 142 |
| The Bottom Line..... | 143 |

Chapter 6 • Understanding the Exchange Autodiscover Process.....145

| | |
|--|-----|
| Autodiscover Concepts | 145 |
| What Autodiscover Provides | 146 |
| How Autodiscover Works | 149 |
| Site Affinity (aka Site Scope) | 157 |
| Planning Certificates for Autodiscover | 158 |
| The X.509 Certificate Standard | 159 |
| Deploying Exchange Certificates..... | 161 |
| The Bottom Line..... | 167 |

Part 2 • Getting Exchange Server Running..... 169

Chapter 7 • Exchange Server 2016 Quick Start Guide171

| | |
|--|-----|
| Server Sizing Quick Reference..... | 172 |
| Hardware..... | 172 |
| Operating Systems | 175 |
| Configuring Windows..... | 178 |
| Active Directory Requirements | 178 |
| Operating System Prerequisites | 179 |
| Installing Exchange Server 2016 | 180 |
| GUI-Based Installation for Mailbox Server Role | 181 |
| Command-Line Installation for Mailbox Server Role..... | 184 |
| Command-Line Installation for Edge Transport Server Role | 185 |
| Post-installation Configuration Steps..... | 186 |
| Final Configuration | 187 |
| Configuring Recipients | 191 |
| The Bottom Line..... | 193 |

| | |
|--|------------|
| Chapter 8 • Understanding Server Roles and Configurations | 195 |
| The <i>Roles</i> of Server Roles | 195 |
| Exchange Server 2016 Server Roles | 197 |
| Mailbox Server | 197 |
| Client Access Services | 202 |
| Edge Transport Server | 204 |
| Possible Role Configurations | 205 |
| Number of Mailbox Servers Deployed | 206 |
| Scaling Exchange Server 2016 Roles | 207 |
| The Bottom Line | 208 |
| Chapter 9 • Exchange Server 2016 Requirements | 209 |
| Getting the Right Server Hardware | 209 |
| The Typical User | 211 |
| CPU Recommendations | 211 |
| Memory Recommendations | 214 |
| Network Requirements | 215 |
| Disk Requirements | 216 |
| Software Requirements | 221 |
| Operating System Requirements | 221 |
| Windows 10 and Windows 8.1 Management Consoles | 223 |
| Additional Requirements | 223 |
| Active Directory Requirements | 224 |
| Installation and Preparation Permissions | 224 |
| Coexisting with Previous Versions of Exchange Server | 225 |
| The Bottom Line | 225 |
| Chapter 10 • Installing Exchange Server 2016 | 227 |
| Before You Begin | 227 |
| Preparing for Exchange 2016 | 228 |
| Existing Exchange Organizations | 229 |
| Preparing the Schema | 229 |
| Preparing the Active Directory Forest | 231 |
| Preparing Additional Domains | 234 |
| Verifying Successful Preparation | 234 |
| Graphical User Interface Setup | 235 |
| Command-Line Setup | 239 |
| Command-Line Installation Options | 240 |
| Command-Line Server-Recovery Options | 242 |
| Command-Line Delegated Server Installation | 243 |
| Installing Language Packs | 244 |
| Removing Exchange Server | 245 |
| The Bottom Line | 246 |

Chapter 11 • Upgrades and Migrations to Exchange Server**2016 or Office 365249**

| | |
|--|-----|
| Upgrades, Migrations, Cross-Forest Migrations, and Deployments | 249 |
| Factors to Consider before Upgrading | 250 |
| Prerequisites | 250 |
| Choosing Your Strategy | 254 |
| Cross-Forest Migration | 255 |
| Upgrading Your Exchange Organization | 257 |
| Office 365 | 260 |
| Office 365 Options | 260 |
| Office 365 Coexistence | 261 |
| Performing a Cross-Forest Migration | 262 |
| Is Cross-Forest Migration the Right Approach? | 262 |
| Choosing the Right Tools | 263 |
| Maintaining Interoperability | 264 |
| Preparing for Migration | 266 |
| Moving Mailboxes | 267 |
| Migrating User Accounts | 268 |
| Permissions Required | 268 |
| Importing Data from PSTs | 268 |
| Tasks Required Prior to Removing Legacy Exchange Servers | 270 |
| Exchange Server Deployment Assistant | 270 |
| The Bottom Line | 271 |

Part 3 • Recipient Administration 273**Chapter 12 • Management Permissions and****Role-Based Access Control275**

| | |
|---|-----|
| RBAC Basics | 275 |
| Differences from Previous Exchange Versions | 275 |
| How RBAC Works | 276 |
| Managing RBAC | 279 |
| Exchange Administration Center | 279 |
| Exchange Management Shell | 280 |
| RBAC Manager | 283 |
| Defining Roles | 284 |
| What's in a Role? | 284 |
| Choosing a Role | 288 |
| Customizing Roles | 289 |
| Distributing Roles | 293 |
| Determining Where Roles Will Be Applied | 293 |
| Assigning Roles to Administrators | 299 |
| Assigning Roles to End Users | 305 |

| | |
|---|-----|
| Auditing RBAC | 310 |
| Seeing What Changes Were Made | 310 |
| Seeing Who Has Been Assigned Rights | 312 |
| The Bottom Line | 313 |

Chapter 13 • Basics of Recipient Management315

| | |
|---|-----|
| Understanding Exchange Recipients | 315 |
| User Mailboxes | 315 |
| Resource and Shared Mailboxes | 316 |
| Mail Users and Mail Contacts | 316 |
| Linked and Remote Mailboxes | 317 |
| Site Mailboxes | 318 |
| Mail-Enabled Groups | 319 |
| Mail-Enabled Public Folders | 319 |
| Defining Email Addresses | 320 |
| Accepted Domains | 321 |
| Email Address Policies | 323 |
| The Bottom Line | 329 |

Chapter 14 • Managing Mailboxes and Mailbox Content331

| | |
|--|-----|
| Managing Mailboxes | 331 |
| Enabling a Mailbox Using the EAC | 332 |
| Enabling a Mailbox Using the EMS | 335 |
| Creating a Mailbox Using the EAC and EMS | 341 |
| Managing User Mailbox Properties | 343 |
| Moving Mailboxes | 356 |
| Moving Mailboxes Using the EAC | 358 |
| Moving Mailboxes Using the EMS | 361 |
| Retrieving Mailbox Statistics | 366 |
| Deleting Mailboxes | 369 |
| Deleting the Mailbox but Not the User | 369 |
| Deleting Both the User and the Mailbox | 370 |
| Purging the Mailbox | 370 |
| Reconnecting a Deleted Mailbox | 371 |
| Bulk Manipulation of Mailboxes Using the EMS | 373 |
| Managing Mailbox Properties Using the EMS | 373 |
| Scripting Account Creation | 375 |
| Managing Mailbox Content | 375 |
| Understanding the Basics of Messaging Records Management | 376 |
| Getting Started with Messaging Records Management | 378 |
| Managing Default Folders | 378 |
| Creating Retention Tags | 379 |
| Managing Retention Policies | 383 |
| The Bottom Line | 386 |

| | |
|---|----------------|
| Chapter 15 • Managing Mail-Enabled Groups, Mail Users, and Mail Contacts | 389 |
| Understanding Mail-Enabled Groups | 389 |
| Naming Mail-Enabled Groups | 390 |
| Creating Mail-Enabled Groups | 391 |
| Managing Mail-Enabled Groups | 395 |
| Creating and Managing Mail Contacts and Mail Users | 405 |
| Managing Mail Contacts and Mail Users Using the EAC | 407 |
| Managing Mail Contacts and Mail Users Using the EMS | 409 |
| The Bottom Line | 410 |
| Chapter 16 • Managing Resource Mailboxes | 413 |
| The Unique Nature of Resource Mailboxes | 413 |
| Exchange 2016 Resource Mailbox Features | 414 |
| Creating Resource Mailboxes | 414 |
| Creating and Configuring Resource Mailboxes | 415 |
| Configuring Advanced Resource Mailbox Features | 417 |
| Configuring Resource Scheduling Policies | 420 |
| Automatic Processing: AutoUpdate vs. AutoAccept | 428 |
| Configuring Resource Calendar Permissions | 429 |
| Creating Room Lists | 430 |
| Converting Resource Mailboxes | 432 |
| The Bottom Line | 433 |
| Chapter 17 • Managing Modern Public Folders | 435 |
| Understanding Architectural Changes for Modern Public Folders | 435 |
| Modern Public Folders and Replication | 436 |
| Modern Public Folder Limitations and Considerations | 437 |
| Moving Public Folders to Exchange Server 2016 | 437 |
| Managing Public Folder Mailboxes | 440 |
| Managing Public Folders | 443 |
| Defining Public Folder Administrators | 449 |
| Using the Exchange Management Shell to Manage Public Folders | 450 |
| Using Outlook to Create a Public Folder | 454 |
| Understanding the Public Folder Hierarchy | 456 |
| Exploring Public Folder High Availability | 457 |
| Managing Public Folder Permissions | 458 |
| Comparing Public Folders, Site Mailboxes, and Shared Mailboxes | 460 |
| The Bottom Line | 460 |
| Chapter 18 • Managing Archiving and Compliance | 463 |
| Introduction to Archiving | 463 |
| Benefits of Archiving | 464 |

| | |
|--|-----|
| Retention | 464 |
| Discovery | 466 |
| Eliminating PST Files..... | 466 |
| Reducing Storage Size | 466 |
| Disaster Recovery | 467 |
| Compliance | 467 |
| Industry Best Practices | 469 |
| Storage Management | 470 |
| Archiving PSTs | 471 |
| Retention Policies | 473 |
| Archiving with Exchange Server 2016 | 474 |
| Exchange In-Place Archive vs. Third-Party Enterprise Archives..... | 474 |
| Retention Policies and Tags..... | 475 |
| Enabling In-Place Archiving | 480 |
| Using the Exchange Server 2016 In-Place Archive | 482 |
| Archive Quotas | 482 |
| Offline Access | 483 |
| Understanding Litigation and In-Place Hold | 483 |
| Placing a Mailbox on In-Place Hold..... | 483 |
| Implementing eDiscovery..... | 484 |
| Requirements and Considerations | 489 |
| Licensing | 489 |
| Server Storage | 489 |
| Client Requirements..... | 490 |
| The Bottom Line..... | 490 |

Part 4 • Server Administration..... 493

| | |
|--|------------|
| Chapter 19 • Creating and Managing Mailbox Databases..... | 495 |
| Getting to Know Exchange Server Databases..... | 495 |
| Exchange Server 2010..... | 495 |
| Exchange Server 2013..... | 496 |
| Exchange Server 2016..... | 496 |
| Basics of Storage Terminology | 496 |
| Storage in Exchange Server 2016 | 498 |
| An Additional Factor: Archive Mailboxes | 499 |
| Disk Size vs. I/O Capacity..... | 500 |
| What's Keeping Me Up at Night?..... | 501 |
| Planning Mailbox Database Storage..... | 501 |
| Maximum Database Sizes | 502 |
| Determining the Number of Databases | 502 |
| Allocating Disk Drives..... | 503 |
| Managing Mailbox Databases | 505 |
| Viewing Mailbox Databases | 506 |
| Creating Mailbox Databases..... | 506 |
| Moving the Mailbox Database EDB File | 507 |

| | |
|---|-----|
| Moving the Mailbox Database Log Files | 507 |
| Properties of a Mailbox Database | 508 |
| The Bottom Line | 517 |

Chapter 20 • Creating and Managing Database

Availability Groups519

| | |
|--|-----|
| Understanding Database Replication in Exchange | |
| Server 2016 | 520 |
| File Mode vs. Block Mode | 520 |
| The Anatomy of a Database Availability Group | 521 |
| File-Share Witness | 524 |
| Creating a Database Availability Group Using the EAC | 525 |
| Creating a Database Availability Group Using EMS | 527 |
| Managing a Database Availability Group | 527 |
| Managing a DAG in EMS | 527 |
| Multiple Databases per Volume | 528 |
| Managing a DAG Network | 530 |
| Adding a Mailbox Database to a DAG | 532 |
| Automatic Reseed (aka AutoReseed) | 539 |
| Understanding Active Manager | 540 |
| Active Manager at Work | 542 |
| Understanding the Best Copy and Server Selection Process | 543 |
| Examples of Best Copy and Server Selection | 547 |
| DAG and Database Maintenance | 549 |
| Incremental Resync | 549 |
| Page Patching | 550 |
| Database Divergence Detection | 550 |
| Applying Updates | 551 |
| Database Schema Updates | 551 |
| Understanding Site Resiliency for Exchange | |
| Server 2016 | 551 |
| DAG Preferred Architecture | 552 |
| DAG for Disaster Recovery | 554 |
| Multiple DAGs | 554 |
| Datacenter Activation Coordination | 555 |
| The Bottom Line | 556 |

Chapter 21 • Understanding the Client Access Services559

| | |
|---|-----|
| Client Access Services Overview | 559 |
| Client Access Services Architecture | 560 |
| Front-End Transport Services | 561 |
| Unified Messaging | 561 |
| Namespace Planning | 562 |
| Namespaces Within a Site | 563 |
| Namespaces Across Multiple Sites | 563 |
| Internal and External URLs | 565 |

| | |
|---|-----|
| Load Balancing | 565 |
| Certificates | 568 |
| Connectivity for Outlook Clients | 573 |
| Autodiscover | 573 |
| Outlook Anywhere | 576 |
| MAPI over HTTP | 578 |
| Exchange Web Services | 578 |
| Modern Attachments | 579 |
| Connectivity for Non-Outlook Clients | 580 |
| Outlook on the Web | 580 |
| ActiveSync | 584 |
| POP3/IMAP | 587 |
| Sharing Between Organizations | 588 |
| Federation Trust | 588 |
| Organization Relationships | 589 |
| Sharing Policies | 590 |
| Securing External Access | 592 |
| Coexisting with Previous Exchange Server Versions | 593 |
| The Bottom Line | 595 |

Chapter 22 • Managing Connectivity with Transport Services597

| | |
|--|-----|
| Understanding the Transport Improvements in Exchange Server 2016 | 597 |
| Message Routing in the Organization | 598 |
| Sending and Receiving Email | 601 |
| Receiving Email from the Internet | 602 |
| Receive Connectors | 603 |
| Sending Email from Your Servers | 610 |
| Send Connectors | 612 |
| Securing Mail Flow | 615 |
| Accepted Domains | 616 |
| Remote Domains | 618 |
| Messages in Flight | 620 |
| Understanding Shadow Redundancy | 620 |
| Understanding Safety Net | 621 |
| Using Exchange Server 2016 Antispam/Anti-Malware Tools | 622 |
| Updating Anti-Malware Engines | 623 |
| Anti-Malware Policy | 623 |
| Managing Anti-Malware Protection | 625 |
| Enabling Antispam Agents | 626 |
| Connection Filtering | 626 |
| Content Filtering | 627 |
| Recipient Filtering | 629 |
| Tarpitting | 630 |
| Sender Filtering | 630 |
| Sender Reputation | 632 |

| | |
|------------------------------------|-----|
| Troubleshooting Email Routing..... | 633 |
| The Bottom Line..... | 634 |

Chapter 23 • Managing Transport, Data Loss Prevention, and Journaling Rules.....635

| | |
|--|-----|
| Introducing the Exchange 2016 Transport Architecture | 635 |
| All Messages Pass Through the Mailbox Server | 636 |
| Setting Up Message Classifications..... | 637 |
| Modifying and Creating Message Classifications..... | 639 |
| Deploying Message Classifications | 640 |
| Setting Up Mail Flow (Transport) Rules..... | 642 |
| Transport Rules Coexistence..... | 643 |
| Transport Rules and Server Design Decisions..... | 644 |
| Selecting Conditions and Exceptions | 645 |
| Creating New Rules with the Exchange Admin Center..... | 646 |
| Creating New Rules with the Exchange Management Shell..... | 649 |
| Introducing Data Loss Prevention..... | 650 |
| Understanding DLP Policies..... | 651 |
| Creating DLP Policies..... | 656 |
| Introducing Journaling | 660 |
| Implementing Journaling | 660 |
| Reading Journal Reports | 662 |
| The Bottom Line..... | 663 |

Part 5 • Troubleshooting and Operating..... 665

Chapter 24 • Troubleshooting Exchange Server 2016.....667

| | |
|---|-----|
| Basic Troubleshooting Principles..... | 667 |
| General Server Troubleshooting Tools | 669 |
| Event Viewer (Diagnostic Logging)..... | 669 |
| The <i>Test-*</i> Cmdlets..... | 672 |
| Troubleshooting Mailbox Servers | 674 |
| General Mailbox Server Health | 674 |
| Using <i>Test-MapiConnectivity</i> | 674 |
| Checking Poison Mailboxes | 675 |
| Checking Database Replication Health..... | 676 |
| Troubleshooting Mail Flow | 677 |
| Using <i>Test-Mailflow</i> | 678 |
| Utilizing the Queue Viewer | 679 |
| Using Message Tracking | 681 |
| Exploring Other Tools | 684 |
| Troubleshooting Client Connectivity | 685 |
| Troubleshooting Autodiscover | 685 |
| Using the <i>Test-* Connectivity</i> Cmdlets..... | 688 |
| The Bottom Line..... | 688 |

| | |
|--|----------------|
| Chapter 25 • Backing Up and Restoring Exchange Server | 691 |
| Backing Up Exchange Server | 691 |
| Determining Your Strategy | 693 |
| Preparing to Back Up and Recover the Exchange Server | 697 |
| Using Windows Server Backup to Back Up the Exchange Server | 697 |
| Performing the Backup | 698 |
| Using Windows Server Backup to Recover the Data | 702 |
| Recovering the Database | 702 |
| Recover Exchange Server Data Using Alternative Methods | 706 |
| Working with Disconnected Mailboxes | 706 |
| Using a Recovery Database | 707 |
| Recovering Single Messages | 709 |
| Recovering Public Folders | 715 |
| Recovering the Entire Exchange Server | 716 |
| Mailbox Server Role | 716 |
| Database Availability Group Members | 717 |
| The Bottom Line | 718 |
| Appendix • The Bottom Line | 719 |
| <i>Index</i> | <i>753</i> |

Introduction

Thank you for purchasing (or considering the purchase of) *Mastering Exchange Server 2016*; this is the latest in a series of Mastering Exchange Server books that have helped thousands of readers to better understand Microsoft's excellent messaging system. Along the way, we hope that this series of books has made you a better administrator and allowed you to support your organizations to the best of your abilities.

When we started planning the outline of this book more than a year before its release, Exchange Server 2016 appeared to be simply a minor series of improvements over Exchange Server 2013. Of course, the further we explored the product, the more we found that was not the case. Many of the improvements in Exchange Server 2016 were major improvements (such as Outlook on the web) and sometimes even complete rewrites (such as in the case of the Client Access services role) of how the product worked previously.

Another challenge then presented itself. The market penetration of Exchange Server 2013 was fairly dominant, but we found that many organizations still run Exchange Server 2010. Therefore, we needed to explain the differences for not only Exchange Server 2013 administrators but also for the Exchange Server 2010 administrators. On the other hand, Exchange Server 2003 reached end-of-life on April 8, 2014. As a result, Microsoft no longer provides security updates, offers free or paid support options, nor provides updated online content such as KB articles for Exchange Server 2003. Organizations with Exchange Server 2003 deployed after April 8, 2014, are responsible for their own support of the product and accept the risk associated with the deployment.

We took a step back and looked at the previous editions of the book to figure out how much of the previous material was still relevant. Some of the material from the Exchange Server 2013 book is still relevant but needed updating. Some required completely rewriting chapters to cover new technologies introduced in Exchange Server 2016 or technologies that have since taken on more importance in deployments and management. We faced the challenge of explaining two management interfaces, Exchange Management Shell and Exchange Admin Center, as well as describing the new roles and features.

We started working with the Exchange Server 2016 code more than a year before we expected to release the book. Much of the book was written using the RTM code that was first made available in October 2015, but as we continued writing the book, we made updates based on changes introduced in Cumulative Update 1 (March 2016). So, you can safely assume when reading this

book that it is based on the latest bits of Exchange Server 2016 that released in late summer 2016. In writing this book, we had a few goals for the book and the knowledge we wanted to impart to the reader:

- ◆ We wanted to provide an appropriate context for the role of messaging services in an organization, outlining the primary skills required by an Exchange Server administrator.
- ◆ We wanted the reader to feel comfortable when approaching an Exchange Server environment of any size. The content in this book can assist administrators of small companies with only one server, as well as administrators who handle large Exchange Server farms.
- ◆ We wanted the skills and tasks covered in this book to be applicable to 80 percent of all organizations running Exchange Server.
- ◆ We wanted the book to educate not only “new to product” administrators but also those “new to version” administrators who are upgrading from a previous version.
- ◆ We wanted the book to familiarize administrators with Office 365 environments and the implementation of hybrid coexistence with on-premises Exchange Server deployments.
- ◆ We wanted to provide familiar references for administrators of previous versions, ensuring that Exchange Server 2010 and 2013 administrators can easily find equivalent solutions in Exchange Server 2016.

Microsoft listened to the advice of many of its customers, its internal consultants at Microsoft Consulting Services (MCS), Microsoft Certified Systems Engineers (MCSEs), Most Valuable Professionals (MVPs), Microsoft Certified Solutions Masters (MCSMs), and Microsoft Certified Trainers (MCTs) to find out what was missing from earlier versions of the product and what organizations’ needs were. Much of this work started even before Exchange Server 2016 was released.

Major Changes in Exchange Server 2016

This book covers the many changes in Exchange Server 2016 in detail, but we thought we would give you a little sample of what is to come in the chapters. As you can imagine, the changes are once again significant, considering the tremendous effort that Microsoft sinks into the Exchange Server line of products. Exchange Server is a significant generator of revenue for Microsoft and is also a foundational service for Office 365. Microsoft has every reason to continue improving this most impressive market leader of email and collaboration services.

The primary changes in Exchange Server 2016 since the latest release (Exchange Server 2013) have come in the following areas:

- ◆ Client access services have been integrated into the Mailbox server role, and the Client Access server role has been removed.
- ◆ Outlook Web App is now known as Outlook on the web, is optimized for tablets, and provides platform-specific experiences for smart phones.

- ◆ MAPI over HTTP is now the default protocol that Outlook uses to communicate with Exchange, which allows a higher level of visibility of transport errors and enhanced recoverability.
- ◆ With SharePoint Server 2016, you can enable Outlook on the web users to link to and share documents stored in OneDrive for Business in an on-premises SharePoint server instead of attaching a file to the message.
- ◆ The Hybrid Configuration Wizard (HCW) is provided as a download to support changes in the Office 365 service and to provide a more stable deployment and consistent experience.
- ◆ Significant enhancements for Data Loss Prevention (DLP) have been added. With a DLP policy and mail flow rules, you can identify, monitor, and protect 80 different types of sensitive information.
- ◆ Public folder integration into the In-Place eDiscovery and Hold workflow enable you to search public folders in your organization and configure an In-Place Hold on public folders.
- ◆ A new eDiscovery search tool, called Compliance Search, provides improved scaling and performance capabilities so you can search very large numbers of mailboxes in a single search.

Of course, many more changes have been introduced in Exchange Server 2016, but the preceding list stands out to us as the most noteworthy improvements. Chapter 2, “Introducing the Changes in Exchange Server 2016,” contains an exhaustive list of all significant changes, as well as changes since specific versions of Exchange Server (for example, Exchange Server 2010 and Exchange Server 2013).

How This Book Is Organized

This book consists of 25 chapters, divided into five broad parts. As you proceed through the book, you’ll move from general concepts to increasingly detailed descriptions of hands-on implementation.

This book won’t work well for practitioners of the time-worn ritual of chapter hopping. Although some readers may benefit from reading one or two chapters, we recommend that you read most of the book in order. Even if you have experience as an Exchange Server administrator, we recommend that you do not skip any chapter, because they all provide new information since the previous iterations of Exchange Server. Only if you already have considerable experience with these products should you jump to the chapter that discusses in detail the information for which you are looking.

If you are like most administrators, though, you like to get your hands on the software and actually see things working. Having a working system also helps many people as they read a book or learn about a new piece of software because this lets them test new skills as they learn them. If this sounds like you, then start with Chapter 7, “Exchange Server 2016 Quick

Start Guide.” This chapter will take you briefly through some of the things you need to know to get Exchange Server running, but not in a lot of detail. As long as you’re not planning to put your quickie server into production immediately, there should be no harm done. Before you put it into production, though, we strongly suggest that you explore other parts of this book. Following is a guide to what’s in each chapter.

Part 1: Exchange Fundamentals

This part of the book focuses on concepts and features of Microsoft’s Windows Server 2012 R2, Exchange Server 2016, and some of the fundamentals of operating a modern client/server email system.

Chapter 1, “Putting Exchange Server 2016 in Context,” is for those administrators who have been handed an Exchange Server organization but who have never managed a previous version of Exchange Server or even another mail system. This will give you some of the basic information and background to help you get started managing Exchange Server and, hopefully, provide a little history and perspective.

Chapter 2, “Introducing the Changes in Exchange Server 2016,” introduces the new features of Exchange Server 2016 as contrasted with previous versions.

Chapter 3, “Understanding Availability, Recovery, and Compliance,” helps even experienced administrators navigate some of the new hurdles that Exchange Server administrators must overcome, including providing better system availability, site resiliency, backup and restoration plans, and legal compliance. This chapter does *not* cover database availability groups in detail; instead, that information is covered in Chapter 20, “Creating and Managing Database Availability Groups.”

Chapter 4, “Virtualizing Exchange Server 2016,” helps you decide whether you should virtualize some percentage of your servers, as many organizations are doing.

Chapter 5, “Introduction to PowerShell and the Exchange Management Shell,” focuses on and uses examples of features that are enabled in PowerShell through the Exchange Server 2016 management extensions for PowerShell. All administrators should have at least a basic familiarity with the Exchange Management Shell extensions for PowerShell even if you rarely use them.

Chapter 6, “Understanding the Exchange Autodiscover Process,” helps you to come up to speed on the inner workings of the magic voodoo that is Autodiscover, a feature that greatly simplifies the configuration of both internal and external clients.

Part 2: Getting Exchange Server Running

This section of the book is devoted to topics related to meeting the prerequisites for Exchange Server and getting Exchange Server installed correctly the first time. While installing Exchange Server correctly is not rocket science, getting everything right the first time will greatly simplify your deployment.

Chapter 7, “Exchange Server 2016 Quick Start Guide,” is where everyone likes to jump right in and install the software. This chapter will help you quickly get a single server up and

running for your test and lab environment. While you should not deploy an entire enterprise based on the content of this one chapter, it will help you get started quickly.

Chapter 8, “Understanding Server Roles and Configurations,” covers the primary services that run on the Exchange Server: mailbox services, transport services, and client access services.

Chapter 9, “Exchange Server 2016 Requirements,” guides you through the requirements (pertaining to Windows Server, Active Directory, and previous versions of Exchange Server) that you must meet in order to successfully deploy Exchange Server 2016.

Chapter 10, “Installing Exchange Server 2016,” takes you through both the graphical user interface and the command-line setup for installing Exchange Server 2016.

Chapter 11, “Upgrades and Migrations to Exchange Server 2016 or Office 365,” helps you decide on the right migration or transition approach for your organization. It recommends steps to take to upgrade your organization from Exchange Server 2010 or 2013 to Exchange Server 2016 or to Office 365. Also included in this chapter are recommendations for migration phases and hybrid coexistence with Office 365.

Part 3: Recipient Administration

Recipient administration generally ends up being the most time-consuming portion of Exchange Server administration. Recipient administration includes creating and managing mailboxes, managing mail groups, creating and managing contacts, and administering public folders.

Chapter 12, “Management Permissions and Role-Based Access Control,” introduces one of the most powerful features of Exchange Server 2016, Role-Based Access Control, which enables extremely detailed delegation of permissions for all Exchange Server administrative tasks. This feature will be of great value to large organizations.

Chapter 13, “Basics of Recipient Management,” introduces you to some concepts you should consider before you start creating users, including how email addresses are generated and how recipients should be configured.

Chapter 14, “Managing Mailboxes and Mailbox Content,” is at the core of most Exchange Server administrators’ jobs since the mailboxes represent the direct customer (the end user). This chapter introduces the concepts of managing mailboxes, mailbox data (such as personal archives), and mailbox data retention.

Chapter 15, “Managing Mail-Enabled Groups, Mail Users, and Mail Contacts,” covers management of these objects, including creating them, assigning email addresses, securing groups, and allowing for self-service management of groups, and it offers guidelines for creating contacts.

Chapter 16, “Managing Resource Mailboxes,” discusses a key task for most messaging administrators. A resource can be either a room (such as a conference room) or a piece of equipment (such as an overhead projector). Exchange Server 2016 makes it easy to allow users to view the availability of resources and request the use of these resources from within Outlook or Outlook on the web.

Chapter 17, “Managing Modern Public Folders,” introduces you to the new public folder storage and management features in Exchange Server 2016. Although public folders are being deemphasized in many organizations, other organizations still have massive quantities of data stored in them. Microsoft has reinvented public folders in this latest release of Exchange Server.

Chapter 18, “Managing Archiving and Compliance,” covers not only the overall concepts of archiving and how the rest of the industry handles archiving but also the exciting archival and retention features.

Part 4: Server Administration

Although recipient administration is important, administrators must not forget their responsibilities to properly set up the Exchange server and maintain it. This section helps introduce you to the configuration tasks and maintenance necessary for some of the Exchange Server 2016 services as well as safely connecting your organization to the Internet.

Chapter 19, “Creating and Managing Mailbox Databases,” helps familiarize you with the changes in Exchange Server 2016 with respect to mailbox database, storage, and basic sizing requirements. Many exciting changes have been made to support large databases and to allow Exchange Server to scale to support more simultaneous users.

Chapter 20, “Creating and Managing Database Availability Groups,” is a key chapter in this book that will affect all administrators from small to large organizations. Exchange Server 2016 relies heavily on Windows Failover Clustering for its site resilience and high availability functionalities. This chapter covers the implementation and management of high availability solutions.

Chapter 21, “Understanding the Client Access Services,” introduces you to the critical client access services and the related components running on the Mailbox server.

Chapter 22, “Managing Connectivity with Transport Services,” brings you up to speed on the Transport services that run with the mailbox and client access services. This chapter discusses mail flow and the transport pipeline in detail.

Chapter 23, “Managing Transport, Data Loss Prevention, and Journaling Rules,” shows you how to implement a feature set that was first introduced in Exchange Server 2007 but has since been greatly improved: the transport rule feature. This chapter also discusses message journaling and Data Loss Prevention policies.

Part 5: Troubleshooting and Operating

Troubleshooting and keeping a proper eye on your Exchange servers’ health are often neglected tasks. You may not look at your Exchange servers until there is an actual problem. In this part, we discuss some tips and tools that will help you proactively manage your Exchange Server environment, ensuring that you can track down problems as well as restore any potential lost data.

Chapter 24, “Troubleshooting Exchange Server 2016,” introduces you not only to troubleshooting the various components of Exchange Server 2016 but also to good troubleshooting techniques. This chapter also includes a discussion of some of the Exchange Server 2016 built-in tools, such as the Exchange Management Shell test cmdlets and the Remote Connectivity Analyzer.

Chapter 25, “Backing Up and Restoring Exchange Server,” includes discussions on developing a backup plan for your Exchange Server 2016 servers as well as how to implement appropriate backup solutions for Exchange Server configuration, databases, logs, and any other relevant information.

Conventions Used in This Book

We use the code-continuation character on PowerShell commands to indicate that the line of text is part of a previous command line.

Many of the screen captures in this book have been taken from lab and test environments. However, sometimes you will see screen captures that came from an actual working environment. We have obscured any information that would identify those environments.

Any examples that include IP addresses have had the IP addresses changed to private IP addresses even if we are referring to Internet addresses.

Remember, Exchange Server is designed to help your organization do what it does better, more efficiently, and with greater productivity. Have fun, be productive, and prosper!

The Mastering Series

The *Mastering* series from Sybex provides outstanding instruction for readers with intermediate and advanced skills, in the form of top-notch training and development for those already working in their field and clear, serious education for those aspiring to become pros. Every *Mastering* book includes the following:

- ◆ Real-World Scenarios, ranging from case studies to interviews, that show how the tool, technique, or knowledge presented is applied in actual practice
- ◆ Skill-based instruction, with chapters organized around real tasks rather than abstract concepts or subjects
- ◆ Self-review test questions, so you can be certain you’re equipped to do the job right



Part 1

Exchange Fundamentals

- ◆ **Chapter 1: Putting Exchange Server 2016 into Context**
- ◆ **Chapter 2: Introducing the Changes in Exchange Server 2016**
- ◆ **Chapter 3: Understanding Availability, Recovery, and Compliance**
- ◆ **Chapter 4: Virtualizing Exchange Server 2016**
- ◆ **Chapter 5: Introduction to PowerShell and the Exchange Management Shell**
- ◆ **Chapter 6: Understanding the Exchange Autodiscover Process**



Chapter 1

Putting Exchange Server 2016 in Context

Email is one of the most visible services that Information Technology (IT) professionals provide; most organizations have become dependent on “soft” information to run their business. As a result, users have developed an attachment to email that goes beyond the hard value of the information it contains. If there’s a problem with email, it affects users’ confidence in their ability to do their jobs—and their confidence in IT.

Microsoft’s Exchange Server products play a key role in electronic messaging, including email. This chapter is a high-level primer on Exchange Server–based email administration and good administration practices, and it prepares you to put Exchange Server 2016 into the proper context. An experienced email administrator may want to proceed to more technical chapters. However, if you are new to the job or need a refresher, or maybe you just want to put email services back into perspective, this chapter is for you!

IN THIS CHAPTER, YOU WILL LEARN TO:

- ◆ Understand email fundamentals
- ◆ Identify email-administration duties

Email’s Importance

If you’re responsible for electronic messaging in your organization, no one has to tell you about its steadily expanding use—you see evidence every time you check the storage space on your disk drives or need an additional tape to complete the backup of your mail server. This section discusses some aspects of electronic mail and the ever-changing nature of email. Even experienced Exchange Server administrators may want to review this section to better understand how their users and requirements are evolving.

Billions of emails are sent every day (more than 200 billion worldwide, according to research firm The Radicati Group). That’s a lot of email messages, on a lot of servers—many of them Exchange servers.

Sure, sending simple text email and file attachments is the most basic function, but email systems (the client and/or the server) may also perform the following important functions:

- ◆ Act as a personal information manager, providing storage for and access to personal calendars, personal contacts, to-do and task lists, personal journals, and chat histories.
- ◆ Provide the user with a single “point of entry” for multiple types of information, such as voicemail, faxes, and electronic forms.

- ◆ Provide shared calendars, departmental contacts, and other shared information.
- ◆ Provide notifications of workflow processes, such as finance/accounting activities, IT events (server status information), and more.
- ◆ Archive important attachments, text messages, and many other types of information.
- ◆ Allow users to access their “email data” through a variety of means, including clients running on Windows computers, Apple computers, Unix systems, web browsers, mobile phones, and even a regular telephone.
- ◆ Perform records management and enable long-term storage of important information or information that must be archived.
- ◆ Enable near-time communication of sales and support information with vendors and customers.

These are just a few of the types of things that an email system may provide to the end user either via the client interface or as a result of some function running on the server.

How Messaging Servers Work

At the core of any messaging system, you will find a common set of basic functions. These functions may be implemented in different ways depending on the vendor or even the version of the product. Exchange Server has evolved dramatically over the past 20 years, and its current architecture is almost nothing like Exchange Server 4.0 from 1996. Common components of most messaging systems include the following:

- ◆ A message transport system that moves messages from one place to another. Examples include the Simple Mail Transport Protocol (SMTP).
- ◆ A message storage system that stores messages until a user can read or retrieve them. Messages may be stored in a client/server database, a shared file database, or even in individual files.
- ◆ A directory service that allows a user to look up information about the mail system’s users, such as a user’s email address.
- ◆ A client access interface on the server that allows the clients to get to their stored messages. This might include a web interface, a client/server interface, or the Post Office Protocol (POP).
- ◆ The client program that allows users to read their mail, send mail, and access the directory. This may include Outlook, Outlook on the web, and a mobile device such as a Windows phone, an iPhone, or an Android device.

Working in tandem with real-time interactive technologies, electronic messaging systems have already produced a set of imaginative business, entertainment, and educational applications with high payoff potential. All of this action, of course, accelerates the demand for electronic messaging capabilities and services.

Most organizations that deploy an email system usually deploy additional components from their email software vendor or third parties that extend the capabilities of the email system or provide required services. These include the following:

- ◆ Integration with existing phone systems or enterprise voice deployments to pull voice messages into the mailbox
- ◆ Message-hygiene systems that help reduce the likelihood of a malicious or inappropriate message being delivered to a user
- ◆ Backup and recovery, disaster recovery, and business continuity solutions
- ◆ Message archival software to allow for the long-term retention and indexing of email data
- ◆ Electronic forms routing software that may integrate with accounting, order entry, or other line-of-business applications
- ◆ Mail gateways to allow differing mobile devices, such as BlackBerry devices, to access the mail server, along with native access through Exchange ActiveSync
- ◆ Email security systems that improve the security of email data either while being transferred or while sitting in the user's mailbox
- ◆ A link load balancer to balance the load between multiple Internet-facing servers or internal servers

What Is Exchange Server?

In its simplest form, Exchange Server provides the underlying infrastructure necessary to run a messaging system. Exchange Server provides the database to store email data, the transport infrastructure to move the email data from one place to another, and the access points to access email data via a number of different clients.

However, Exchange Server, when used with other clients such as Outlook or Outlook on the web, turns the “mailbox” into a point of storage for personal information management such as your calendar, contacts, task lists, and any file type. Users can share some or all of this information in their own mailbox with other users on the message system and start to collaborate.

The Outlook and Outlook on the web clients also provide access to public folders. Public folders look like regular mail folders in your mailbox, except that they are in an area where they can be shared by all users within the organization. A folder can have specialized forms associated with it to allow the sharing of contacts, calendar entries, or even other specialized forms. Further, each public folder can be secured so that only certain users can view or modify data in that folder.

The Unified Messaging features in Exchange Server 2016 further extend the functions of Exchange Server in your organization by allowing your Exchange Server infrastructure to also act as your voicemail system and direct voicemails and missed-call notifications automatically to the user's mailbox.

While integrated voicemail solutions are nothing new for Exchange Server customers, Microsoft is now providing these capabilities out of the box rather than relying on third-party products.

Exchange Server 2016 tightens the integration of collaborative tools in its integration with Skype for Business Server 2015, the Skype for Business client, and the Skype for Business mobile client. Skype for Business provides a core set of Session Initiation Protocol (SIP)-based enterprise voice capabilities that allows it to act as a PBX in many cases. With Exchange Server, Skype for Business, Outlook, and the Skype for Business client, users enjoy full Unified Messaging with software-based telephony from their computer, including the voicemail and missed-call notification provided by Exchange Server and Outlook. Furthermore, Skype for Business can log chat and instant-message conversation logs to a folder in the user's mailbox. Exchange Server 2016 further pushes this integration, embedding basic instant messaging (IM) and presence capabilities into the Outlook on the web premium experience.

The capabilities of the client can be extended with third-party tools and forms-routing software so that electronic forms can be routed through email to users' desktops.

About Messaging Services

Electronic messaging is far more than email. Together, Exchange Server 2016 and its clients perform a variety of messaging-based functions. These functions include email, unified messaging, message routing, scheduling, and support for several types of custom applications. Together these features are called messaging services.

Many Modes of Access

For years, the only way to access your email system was to use a Windows, Mac, or Unix-based client and access the email system directly. In the case of Outlook and Exchange Server, this access was originally in the form of a MAPI client directly against the Exchange server. As Exchange Server has evolved, it has included support for RPC over HTTP, MAPI over HTTP, Exchange Web Services (EWS), and finally mobile device access (via ActiveSync). Exchange Server 2016 doesn't offer any radically new modes of mailbox access as Exchange Server 2007 did, but it does provide ongoing support and refinement of existing Exchange Server 2007 technologies, such as Exchange Web Services, that can provide additional mechanisms for accessing data in mailboxes and a move away from RPC in client connectivity in favor of Outlook on the web and mobile devices.

Outlook on the web (formerly Outlook Web Access) has evolved quickly and, in Exchange Server 2016, bears almost no resemblance to the original version found in Exchange Server 5.0 in terms of features, functions, and the look of the interface. Exchange Server 2016 Outlook on the web is a step beyond Exchange Server 2013. It expands the previous option configuration experience of the Exchange Control Panel (ECP), which gives users a much greater degree of control over their mailboxes, contacts, and group memberships. ECP is built into the Outlook on the web interface. Using ECP, end users can create and join distribution groups (where permissions have been assigned), track their own messages throughout the organization, and perform other functions that in Exchange 2010 and earlier versions required help-desk or IT professional intervention. Another significant feature of Outlook on the web is the ability to use the web-based interface when working offline and completely disconnected from the network.

With Exchange Server 2016, Exchange ActiveSync (EAS) continues to offer significant partnerships with and control over mobile devices. Many vendors have licensed EAS to provide their mobile devices with a high-performance, full-featured push mobile synchronization experience that extends beyond mobile phones and into tablet devices.

With all of these mechanisms for retrieving and sending email, it is not unusual for users to access their mailboxes using more than one device. In some cases, we have seen a single user accessing her mailbox from her desktop computer, her tablet device using Outlook Anywhere, and her Windows Phone device.

In medium and large organizations, the fact that users are accessing their mailboxes from more than one device or mechanism will affect not only hardware sizing but also, potentially, your licensing costs.

How Messaging Services Are Used

Certainly, email is a key feature of any messaging system, and the Outlook Calendar is far better than previous versions of Microsoft's appointment and meeting-scheduling software. Outlook 2016 together with Exchange Server 2016 introduces even more synergy. Figure 1.1 and Figure 1.2 show the Outlook 2016 client Calendar and Inbox in action.

FIGURE 1.1
Outlook 2016
Appointment
scheduling on an
Exchange Server
2016 mailbox

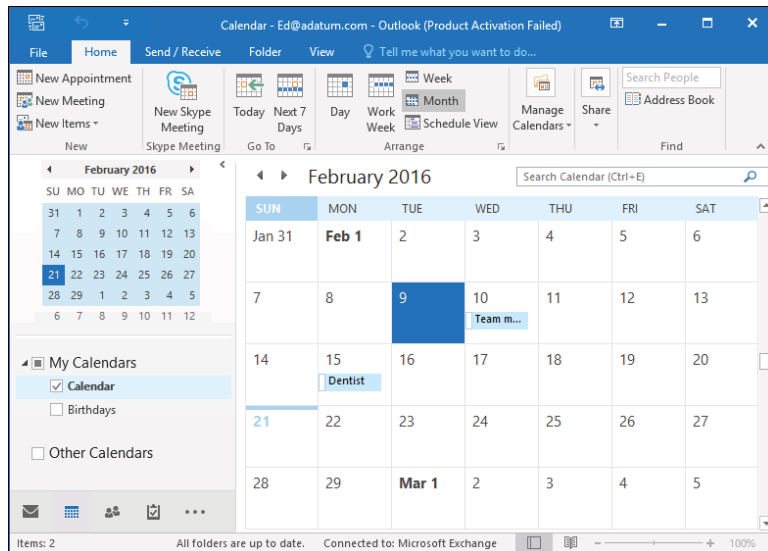


Figure 1.3 shows the new Outlook on the web 2016 web browser client. Outlook on the web provides the full, premium user experience for browsers other than Internet Explorer; it also supports Mac OS X Safari, Firefox, and Chrome. Those coming from older versions of Exchange Server will immediately notice a cleaner, less-cluttered interface and new functionalities such as Offline Usage.

FIGURE 1.2
The Outlook 2016 client Inbox on an Exchange Server 2016 mailbox

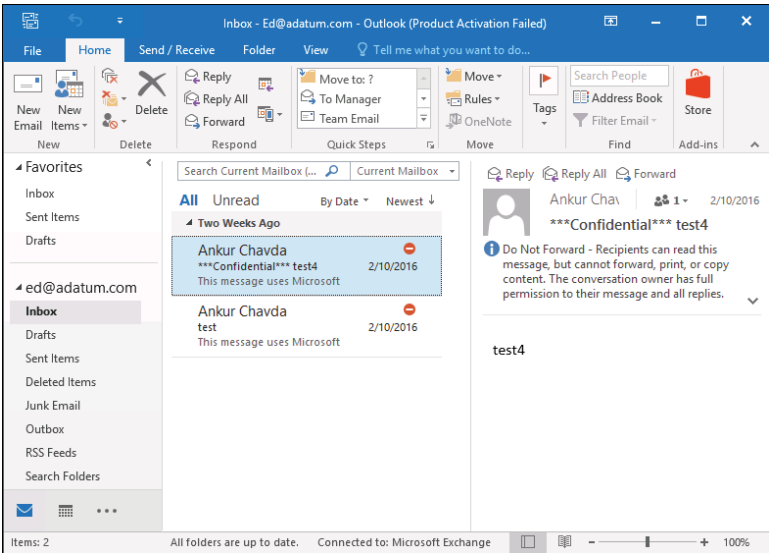
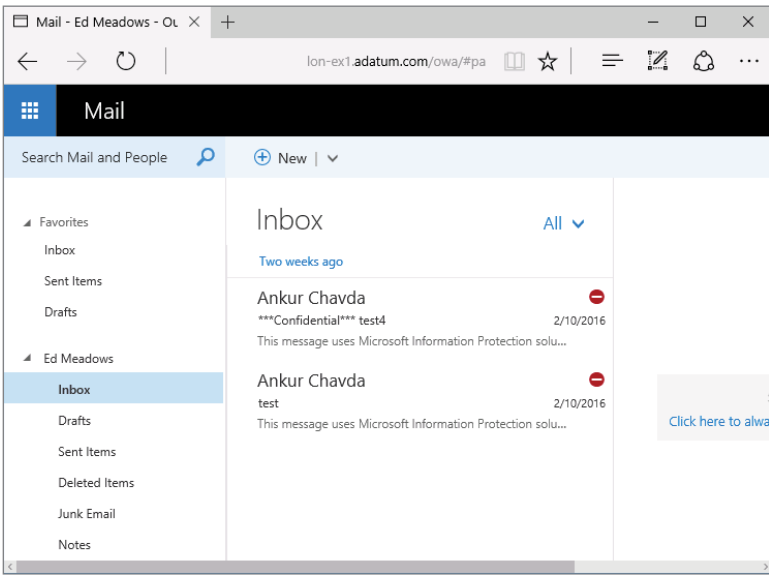


FIGURE 1.3
Outlook on the web on an Exchange Server 2016 mailbox



Email clients are exciting and sexy, but to get the most out of Exchange Server 2016 you need to throw away any preconceptions you have that messaging systems are only for email and scheduling. The really exciting applications are not those that use simple email or scheduling but those that are based on the routing capabilities of messaging systems. These applications bring people and computers together for improved collaboration.

The Universal Inbox

Email systems are converging with their voicemail and enterprise voice-solution cousins. The concept of unified messaging is nothing new to email users. For the past 20 years, third-party

vendors have included email integration tools for voicemail, network faxing solutions, and third-party integration. However, for most organizations, integrated voicemail remains the exception rather than the rule. Exchange Server 2007 introduced integrated voice, which Exchange Server 2016 continues to improve.

Organizations with IP-based telephone systems or telephone systems with an IP gateway can easily integrate a user's voicemail with the Exchange Server user's mailbox. The Exchange Server 2016 Unified Messaging features handle the interaction between an organization's telephone system and Exchange Server mailboxes. Inbound voicemail is transferred into the user's mailbox as a cross-platform-friendly MP3 file attachment; this message includes an Outlook or Outlook on the web form that allows the user to play the message. As well, the voicemail text can be transcribed into the body of the email message for quick reading by the user during meetings or rapid glancing at the Inbox. Because the default format is MP3 in Exchange Server 2016 (it was a Windows Media file in Exchange Server 2007, using a custom codec), this file can be easily played on mobile devices from any manufacturer, allowing easy on-the-go access to voicemail. A short voicemail message may be anywhere from 40 KB to 75 KB in size, whereas longer voicemail messages may range from 200 KB to 500 KB in size. One estimate that is frequently used for the size of a voicemail message is around 5 KB per second of message.

Inbound voicemail increases the demands on your Exchange server from the perspective of required disk space and possible additional server hardware. As an administrator, you need to consider this.

JUST THE FAX, MA'AM

In Exchange Server 2007, the Unified Messaging features included the out-of-the-box capability to capture incoming facsimile (fax) messages. There were some limitations, but it provided good basic functionality. For outbound fax capability, organizations had to deploy some other solution, typically a third-party fax package.

Since Exchange Server 2010, Microsoft made the decision to cut this feature. When talking with the product group, it's not hard to figure out why; the inbound-only fax functionality wasn't enough for the customers who needed fax integration. Exchange Server needed to either add outgoing fax capability and beef up its feature set (and lose other desired functionality) or drop the existing functionality because the majority of Exchange Server 2007 customers needed a third-party product anyway. Although it's always disappointing to lose a feature, most of the organizations we've talked to didn't use it to begin with. We think that Microsoft definitely made the right call, if you'll pardon the pun.

Architecture and Core Functionality Overview

Understanding a bit about how Exchange Server works from an architectural perspective will help make you a better administrator. You don't have to be able to reproduce or write your own client/server messaging system, but it helps to know the basics.


THE EXTENSIBLE STORAGE ENGINE

The Exchange Server database uses a highly specialized database engine called the Extensible Storage Engine (ESE). Generically, you could say it is almost like SQL Server, but this is technically not true. It is a client/server database and is somewhat relational in nature, but it is

designed to be a single-user database (the Exchange server itself is the only component that directly accesses the data). Further, the database has been highly tuned to store hierarchical data, such as mailboxes, folders, messages, and attachments.

Without going into a lot of techno-babble on the database architecture, it is important that you understand the basics of what the database is doing. Figure 1.4 shows conceptually what is happening with the ESE database as data is sent to the database. In step 1, an Outlook client sends data to the Exchange server (the Information Store service); the Information Store service places this data in memory and then immediately writes the data out to the transaction log files associated with that database.

FIGURE 1.4
Exchange data and
transaction logs



The transaction log that is always written to is the current transaction log for that particular database (e00.log, for example). Each transaction log file is exactly 1 MB in size, so when the transaction log is filled up, it is renamed to the next sequential number. For example, an old transaction log file might be named like this: e000004032.log. We often get questions about the logic of the transaction logs, and how they reserve space on the disk, whether they are empty or full. An easy way to look at it is to compare a log file to a carton of milk. When you have a carton of milk, it always takes up the same space in your fridge, empty or full. The same is true of the log files. Empty log files (current log file and *reserved* log files) are empty, or partially full; the renamed, *old*, log files are full. However, they take up the same amount of space on the disk.

The data, such as new email messages that enter the organization, is retained in RAM for some period of time (maybe as little as 5 seconds or maybe even 60 seconds or more) before it is flushed to the database file. The actual period that data is retained in memory will depend on how much cache memory is available, what types of operations are happening in the data, and how busy the server is. The important operation, though, is to make sure that as soon as the data is sent to the Exchange server, it is immediately flushed to the transaction log files. If the server crashes before the data is written to the database file, the database engine (the store process) will automatically read the transaction log files once the server is brought back up and compare them to the data that's stored in the corresponding mailbox databases. Any inconsistency is resolved by replaying the missing data operations from the transaction logs back into the database, assuming that the entire transaction is present; if it's not, the operations are not written (and you can be confident that the operation wasn't completed at the time the crash happened). This helps ensure that the integrity of the mailbox database is preserved and that half-completed data operations aren't written back into the database and allowed to corrupt good data.

The transaction log files are important for a number of reasons. They are used by Microsoft replication technologies (as you'll learn in Chapter 19, "Creating and Managing Mailbox Databases"), but they can also be used in disaster recovery. The transaction logs are not purged

off the log disk until a full backup is run; therefore, every transaction that occurred to a database (new data, modifications, moves, deletes) is stored in the logs. If you restore the last good backup to the server, Exchange Server can replay and rebuild all the missing transactions back into the database—provided you have all the transactions since the last full backup.


In early versions of Exchange Server, you had two separate mail store objects: the *storage group*, which was a logical container that held an associated set of transaction logs, and the *mailbox database*, a set of files that held the actual permanent copies of user mailboxes. You often had multiple mailbox databases per storage group, meaning that one set of transaction logs contained interwoven transaction data for multiple databases (which could have detrimental effects on performance, space, and backups).

In Exchange Server 2016, you still have mailbox databases. However, storage groups were removed in Exchange Server 2010; each mailbox database now has its own integral set of transaction log files. In fact, mailbox databases—which were once tightly coupled with specific servers—can have copies on multiple servers in the organization, even spread across multiple sites. This functionality was introduced by moving the mailbox databases from the Server hierarchy to the Organization hierarchy, essentially rendering them a *shared* object that can become active on any server in the organization. The *database availability group* container is now available to contain servers that participate in the replication of mailbox databases with each other.

EXCHANGE AND ACTIVE DIRECTORY

We could easily write two or three chapters on how Exchange Server interacts with Active Directory, but the basics will have to do for now. Exchange Server relies on Active Directory for information about its own configuration, user authentication, and email-specific properties for mail-enabled objects such as users, contacts, groups, and public folders. Look at Figure 1.5 to see some of the different types of interactions that occur between Exchange Server and Active Directory.

FIGURE 1.5
Active Directory
and Exchange
Server



Because most of the Exchange Server configuration data for an Exchange server is stored in Active Directory, all Exchange Server roles must contact a domain controller to request its configuration data; this information is stored in a special partition of Active Directory database called the *configuration partition*. The configuration partition is replicated to all domain

controllers in the entire Active Directory forest. Note that you can have only a single Exchange organization per Active Directory forest.

Each of the Exchange Server components uses Active Directory for different things. Some of those functions include:

Mailbox Components For mailbox operations, Exchange Server must query Active Directory to authenticate users, enumerate permissions on mailboxes, look up individual mailbox limits, and determine which mailboxes are on a particular server. They also require access to global catalog servers to look up email addressing information, distribution list membership information, and other data related to message routing.

Client Access Components For client access, Exchange Server requires access to Active Directory to look up information about users, Exchange ActiveSync, and Outlook on the web user restrictions.

Controlling Mailbox Growth

As users have become more savvy and competent at using Outlook and the features of Exchange Server, and email messages themselves have become more complex, the need for email storage has grown. Back in the days of Exchange Server 4.0, an organization that gave its users a 25 MB mailbox was considered generous. With Exchange Server 2003, a typical user's mailbox may have a storage limit of 300 to 500 MB, with power users and VIPs requiring even more. At TechEd 2006, Exchange Server gurus were tossing about the idea that in the future a default mailbox limit would be closer to 2 GB as users start incorporating Unified Messaging features. Current discussions now look forward to and assume unlimited-sized mailboxes within the next few years.

We all see users with mailbox sizes in the gigabyte range, but is your organization prepared for a typical user with an unlimited mailbox size? What sort of concerns will you face when your average user has 25 GB, 50 GB, 100 GB, or even unlimited content (not just email!) in their mailbox?

Certainly, the need for more disk storage will be the first factor that organizations need to consider. However, disk storage is reasonably cheap, and many larger organizations that are supporting thousands of mailbox users on a single Mailbox server already have more disk space than they can practically use. This is due to the fact that they require more disk spindles to accommodate the number of simultaneous I/Os per second (IOPS) that are required by a large number of users. While early versions of Exchange Server were primarily *performance-bound*—meaning that they would require more drive performance before they required more disk capacity—versions since Exchange Server 2007 have solidly pushed that to being *capacity-bound*. With the performance characteristics and capacities of modern drives, it becomes feasible to economically provision Exchange Server storage in support of large mailboxes.

For most administrators with large amounts of mail storage, the primary concern they face is the ability to quickly and efficiently restore data in the event of a failure. These administrators are often faced with service-level agreements that bind them to maximum restoration times. In even the most optimal circumstances, a 300 GB mailbox database will take some time to restore from backup media. However, these issues have largely been mitigated by the use of database availability groups (DAGs), which ensure constant *copies* of mailbox databases that reside on other servers, essentially providing a constant *live* backup of mailbox databases on other servers, and in other datacenters.

Microsoft recommends that you do not allow an Exchange Server mailbox database to grow larger than 200 GB unless you are implementing continuous-replication technologies in

Exchange Server 2016. If you use database availability groups to replicate databases to multiple servers, the maximum database size recommendation goes up (way up) to 2 TB. However, the maximum supported database size is actually 64 TB. If you require more than the maximum recommend database storage, Exchange Server 2016 Standard Edition allows you to have up to 5 mailbox databases and Exchange Server 2016 Enterprise Edition allows you to have up to 100.

The solution in the past was to restrain the user community by preventing them from keeping all of the mail data that they might require on the mail server. This was done by imposing low mailbox limits, implementing message-archival requirements, keeping deleted items for only a few days, and keeping deleted mailboxes for only a few days.

However, as Unified Messaging data arrives in a user's mailbox and users have additional mechanisms for accessing the data stored in their mailbox, keeping mail data around longer is a demand and a requirement for your user community. The Exchange Server 2016 archive mailbox feature also drives the need for more storage, as message archival moves away from the PST files and back into Exchange Server in the form of archive mailboxes. Those archive mailboxes can be *segregated* to a dedicated mailbox database and be set to a different backup schedule and their own set of management practices.

Personal Folders or PST Files

While we're on the subject of PST files, let's discuss this pesky feature of client management. The Outlook Personal Folder, or PST files, can be the very bane of your existence. Outlook allows users to create a local database, named Personal Folder, in which users can create folders and archive email. Although this seems like a good feature on the surface, there are a few downsides:

- ◆ Once data is in a user's PST file, you, as the server administrator, have lost control of it. If you ever had to find all copies of a certain message, perhaps for a lawsuit, you would be out of luck. PSTs can become a management and security nightmare as data is suddenly distributed all over your network.
- ◆ The data in PST files take up more space than the corresponding data on the server.
- ◆ The default location for a PST is the local portion of the user's profile; this means it is stored on the local hard disk of their computer and is not backed up.
- ◆ PST files can get corrupted, become misplaced, or even be lost entirely. PSTs are not designed for access over a network connection; they're meant to be on the local hard drive, which wastes space, as well as complicates the backup and management scenarios.
- ◆ Starting with Exchange Server 2010, Personal Archives stored on the server can be populated from PST files, therefore offering a true alternative to those pesky local files.

Email Archiving

Sometimes, managing a mail server seems like a constant race between IT and users to keep users from letting their mailbox run out of space. Users are pack rats and generally want to keep everything. If there is a business reason for them to do so, you should look at ways to expand your available storage to accommodate them.

However, as databases become larger and larger, the Exchange server will be more difficult to manage. You might start requiring hundreds and hundreds of gigabytes (or even terabytes) of storage for email databases. Worse still, performing backups and data recovery take longer.

Exchange Server 2016 provides some archiving features, such as the Personal Archive. Also, large mailboxes could be moved to an Office 365 subscription, in a hybrid coexistence model.

For those organizations that are *not* opting to head out to the cloud or do *not* choose Office 365 as their email solution, this is where email archiving becomes useful. The last time we counted, several dozen companies were in the business of supplying email archiving tools and services. Archiving products all have a lot of functions in common, including the ability to keep data long term in email archival, to allow the users to search for their own data, and to allow authorized users to search the entire archive.

If you look at how email is archived, archive systems generally come in one of three flavors:

- ◆ Systems that depend on journaling to automatically forward every email sent or received by specified users on to the archive system.
- ◆ Systems that perform a scheduled “crawl” of specified mailboxes, looking for messages that are eligible to be moved or copied to the archive.
- ◆ Systems that move data to the archive by copying the log files from the production Mailbox servers and then replaying the logs in to the archive. This is called log shipping.

Each of these methods has its advantages and disadvantages with respect to using storage, providing a complete archive, and dealing with performance overhead.

In the previous section, we discussed briefly the archive mailbox as an alternative to the management of PST files. However, its ability goes beyond the manual move of email messages to a dedicated location on the server. For any user who requires email archival, a Personal Archive can be created for that user. As email ages past a certain point, the mail is moved from the active mailbox to the archive mailbox by using Archive Policies. The user can still access and search the archive mailbox from Outlook on the web or Outlook, though. The email data remains on the Exchange server and, therefore, does not require an additional email archival infrastructure.

We often are asked if this information can be made available offline; keep in mind that it cannot. Personal Archives cannot be included in Offline Stores (OST) files. This is by design, and we’re kind of glad that it works this way, because we are continuously trying to reduce the email footprint on the client computers. OST files get very large, very fast, and can cause plenty of headaches as well. Note that with Outlook 2013 and Outlook 2016, you can adjust how many days, weeks, months, or years to sync offline.

IF I USE A THIRD-PARTY SOLUTION, DOES IT MATTER HOW I ARCHIVE?

Every third-party archival vendor is going to tell you how their product is best and give you long technical reasons why their approach is so much better than the competition’s. The dirty little secret is that all three approaches have their pros and cons:

Journaling is based on SMTP. If content doesn’t run across SMTP, it won’t get journaled and, therefore, won’t get archived. Journaling is great for capturing messaging and calendaring traffic that involves multiple parties or external entities, but it won’t capture what happens to messages and other mailbox data once they’re in the mailbox. Journaling can also place an additional load on the Hub Transport servers, depending on the amount and type of messaging traffic your users generate.

Crawling can capture changes only at certain intervals; it can't capture every single change, even though it overcomes many of the limitations of journaling. For example, if one user sends a message to another in violation of policy and both hard-delete their copy of the message before the next crawl interval, that message won't be detected and archived. The more often you schedule the crawl, the more of a performance impact your Mailbox servers will suffer.

Log shipping is the best of all options; it captures every transaction and change, allowing you to capture the entire history of each object while offloading the performance hit from your Exchange servers. However, the Exchange Server product team does not like the concept of log shipping and tries to discourage its use—mainly because there are vendors who try to inject data back into Exchange Server by modifying logs. This, needless to say, results in mailbox data that won't be supported by Microsoft.

Public Folders

The end-user experience for public folders has not changed in Exchange Server 2016, though the architecture has changed in recent years—mainly the storage of the public folders, which is now in a mailbox database, instead of the public folder database. Public folders are for common access to messages and files. Files can be dragged from file-access interfaces, such as File Explorer, and dropped into public folders. The whole concept of public folders has many organizations in a quandary as they try to figure out the best place for these collaborative applications. Increasingly, applications that were once “best suited” for a public folder are now better suited for web pages or portals, such as SharePoint workspaces. Although the whole concept of public folders is perceived as being deemphasized since Exchange Server 2007, Microsoft continues to support public folders, and many organizations will continue to find useful applications for public folders for the foreseeable future.

A key change in public-folder storage occurs in Exchange Server 2016, one that finally breaks the paradigm of dedicated public folder databases and public folder replication. Although we discuss this change in Chapter 2, “Introducing the Changes in Exchange Server 2016,” we just briefly note here that public folders are now stored in mailbox databases and can be replicated as mailbox database copies in a database availability group.

You can set up sorting rules for a public folder so that items in the folder are organized by a range of attributes, such as the name of the sender or creator of the item or the date that the item was placed in the folder. Items in a public folder can be sorted by conversation threads. Public folders can also contain applications built on existing products such as Word or Excel or built with Exchange Server or Outlook Forms Designer, client or server scripting, or the Exchange Server API set. You can use public folders to replace many of the maddening paper-based processes that abound in every organization.

For easy access to items in a public folder, you can use a *folder link*. You can send a link to a folder in a message. When someone navigates to the folder and double-clicks a file, the file opens. Everyone who receives the message works with the same linked attachment, so everyone reads and can modify the same file. As with document routing, applications such as Microsoft Word can keep track of each person's changes to and comments on file contents. Of course, your users will have to learn to live with the fact that only one person can edit an application file at a

time. Most modern end-user applications warn the user when someone else is using the file and if so allow the user to open a read-only copy of the file, which of course can't be edited.

Things Every Email Administrator Should Know

The information in this section is something that we often find even our own email administrators and help-desk personnel unaware of. Sometimes the most important skill any technology administrator has is not a specific knowledge of something but generic knowledge that they can use to quickly find the right answer.

A Day in the Life of the Email Administrator

We know and work with a lot of email administrators, and we can honestly say that no two people have the same set of tasks required of them. Your CEO, director of information technology, or even your supervisor is going to ask you to pull rabbits out of your hat, so don't expect every day to be the same as the last one. (And invest in some rabbits.) Keep up with your technology and supporting products so that you can be ready with answers or at the very least intelligent responses to questions.

DAILY ADMINISTRATIVE TASKS

So, what are some typical tasks that you may perform as part of your duties as an email administrator? These tasks will depend on the size of your organization, the number of administrators you have running your Exchange Server organization, and how administrative tasks are divided up.

Recipient Management Tasks These are certainly the biggest day-to-day tasks that most Exchange Server administrators in medium and large organizations will experience. Recipient management tasks may include:

- ◆ Assigning a mailbox to a user account
- ◆ Creating mail-enabled contacts
- ◆ Creating and managing mail groups
- ◆ Managing mail-enabled object properties such as users' phone numbers, assigning more email addresses to a user, or adding/removing group members

Basic Monitoring Tasks These ensure that your Exchange servers are healthy and functioning properly:

- ◆ Checking queues for stalled messages
- ◆ Verifying that there is sufficient disk space for the databases and logs
- ◆ Making sure that the message-hygiene system is functioning and up-to-date
- ◆ Running and verifying daily backups
- ◆ Reviewing the event logs for unusual activity, errors, or warnings
- ◆ Checking Performance Monitor to gauge how the Exchange servers are performing

Daily Troubleshooting Tasks These include the following:

- ◆ Reviewing nondelivery report messages and figuring out why some mail your users are sending might not have been delivered
- ◆ Looking up errors and warnings that show up in the event logs to determine if they are serious and warrant corrective action
- ◆ Looking at mail flow in the organization to identify why delivery to some recipients is taking a long time

Security-Related Tasks Some of these are performed daily, while others are performed only weekly or monthly:

- ◆ Looking at server and service uptimes to ensure that servers are not rebooting unexpectedly
- ◆ Reviewing the event logs for warnings that may indicate users are inappropriately accessing other users' data
- ◆ Saving the IIS (Internet Information Services) and SMTP and connectivity logs or even reviewing their content

Email Client Administration Tasks These include the following:

- ◆ Troubleshooting Autodiscover connectivity and client issues
- ◆ Diagnosing problems with mobile or tablet devices that use Exchange ActiveSync connectivity

Application Integration Tasks These are performed on an as-needed basis and may include the following:

- ◆ Establishing and diagnosing SMTP connectivity with email-enabled third-party applications such as web servers
- ◆ Configuring, testing, and troubleshooting Unified Messaging interoperability with voice and Session Initiation Protocol (SIP) systems
- ◆ Configuring, testing, and troubleshooting connectivity with SharePoint Server site mailboxes

COMMUNICATING WITH YOUR USERS

Communicating with your users is probably one of the most important things you do. Keeping your users informed and delivering good customer service are almost as important as delivering the IT service itself. Keeping users informed of full or partial service outages such as mobile or iPhone support or web connectivity may not score any immediate points, but users appreciate honest, forthright information. Remember how you felt the last time you were waiting for an airplane to arrive that kept on being delayed and delayed, and all the airline could do was be evasive?

Also, remember to have multiple avenues of communication available to your users. For example, you may need to get out to your users the message that you will be having downtime on the weekend. Postings on your company intranet or even the bulletin board in the cafeteria or on the wall of the elevator are good ways to keep your users informed.

PREPARING REPORTS

Maybe we have just worked in large IT environments for too long now, but it seems to us that information technology is more and more about reports and metrics. We are frequently asked to provide reports, statistics, and information on usage—not necessarily information on performance (how well the system performed for the users) but other types of metrics. Depending on your management, you may be asked to provide the following:

- ◆ Total number of mailboxes and mailbox sizes
- ◆ Top system users and top source/destination domains
- ◆ Antispam and message-hygiene statistics
- ◆ Disk space usage and growth
- ◆ System availability reports indicating how much unscheduled downtime may have been experienced during a certain reporting period
- ◆ Total number of messages sent and received per day
- ◆ Average end-to-end email delivery time

Exchange does not provide you with a way to easily access most of this data. The mailbox statistics can be generated using the Exchange Management Shell, but many of these will actually require an additional reporting product, such as System Center 2012 R2.

Something that you can do to prepare for a reporting requirement is to ensure that you are keeping two to four weeks' worth of message-tracking and protocol logs.

SCHEDULED DOWNTIME, PATCHES, AND SERVICE PACKS

As the discussion over moving to “the cloud” becomes more prevalent in most industries, the common argument that keeps on coming back in favor for moving Exchange Server services to some version of Exchange Online or Office 365 is server availability. No one likes downtime, whether it is scheduled or not. Management may actually be holding you to a specific service-level agreement (SLA) that requires you to provide so many hours of uptime per month or to provide email services during certain hours. Unscheduled downtime is anything that happens during your stated hours of operation that keeps users from accessing their email.

Even a small organization can provide very good availability for its mail services, and without large investments in hardware. Good availability begins with the following:

- ◆ Server hardware should always be from a reputable vendor and listed in the Microsoft Server Catalog.
- ◆ Server hardware should be installed using the vendor recommended procedures and updated regularly. Problems with servers are frequently caused by outdated firmware and device drivers.
- ◆ Once the server is in production, it should not be used as a test bed for other software. Keep an identically configured server that uses the same hardware for testing updates.

Don't underestimate the importance of training and documentation. In general, the industry formula for providing better availability for any system is to spend more money to purchase redundant servers and build failover clusters. But often better training for IT personnel and a simple investment in system documentation, as well as system policies and procedures, can improve availability—and for less money.



Real World Scenario

INTERNAL STAFF TRAINING IS JUST AS IMPORTANT AS YOUR INFRASTRUCTURE

Company LMNO P invested hundreds of thousands of dollars in their infrastructure to improve server uptime. Three months into the operation of the new system, an untrained operator accidentally brought down a 15,000-mailbox database availability group (DAG) simply because he had been asked to do a task he had never done before and the organization did not have documentation on how to proceed. So keep in mind that documentation, training, and procedures are very important in improving uptime.

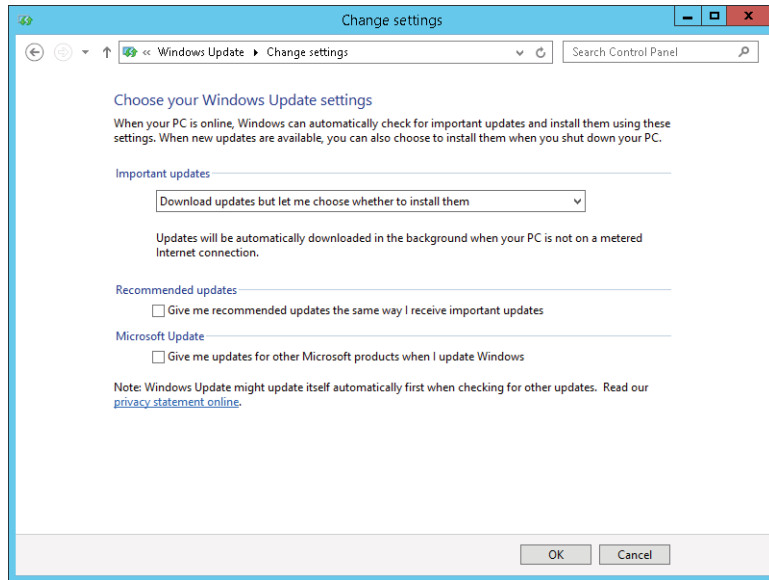
Even the biggest mailbox servers in large database availability groups need some scheduled downtime. Even if it is scheduled in the wee hours of the morning, undoubtedly someone, somewhere, somehow will need access when you are working on the system. Thankfully, the DAG solution for high availability ensures that users may never notice the scheduled server downtime, since mailbox services can be switched over to another member server in the DAG. That being said, when you are driving your car with no spare tire in the trunk, you are more vulnerable to a flat tire. The same is true of the DAG, because when a member server is offline for maintenance, the DAG loses a potential mailbox server that is capable of taking over in the event of server failure.

When your scheduled downtime will affect components that can impact server availability for your users, that downtime should be well communicated. Also, you should document your scheduled downtime as part of your operational plans and let your user community know about these plans. The specific time window for maintenance should always be the same; for some organizations, this might be 6:30 PM to 10:30 PM on Thursday once per month, whereas other organizations might schedule downtime from 11:00 PM Saturday until 4:00 AM every Sunday.

The number-one reason for downtime is to apply updates and fixes to the operating system or to the applications running on the server. Microsoft releases monthly security updates for the operating system and applications if vulnerabilities are discovered. Every few months, Microsoft releases updates for Exchange Server 2016 that fix bugs or that may even add slight functionality. New for Exchange Server 2016, Microsoft uses a quarterly update release cycle. Each quarter, a cumulative update (CU) is released for Exchange Server 2016. You can install the CU in your environment to update it with the latest updates and fixes, and you do not need to install previous CUs before you install the latest CU.

Microsoft's updates are usually downloaded to your servers shortly after they are released. The server can download them directly from Microsoft, or they can be downloaded from Windows Software Update Service (WSUS), Microsoft System Center Configuration Manager 2012 R2, or another third-party server inside your network. Whichever you choose, it is important that you make sure that the machine is a server and not a workstation. For example, make sure the automatic updates component of Windows Server is configured correctly. Figure 1.6 shows the Change Settings options for Windows Update.

FIGURE 1.6
Configuring auto-
matic updates



For production Exchange servers, you should configure the server with the option Download Updates But Let Me Choose Whether To Install Them. This is an important setting because if you choose the Install Updates Automatically (Recommended) option, the server will automatically apply any update within a day or so of downloading it. This is not a desirable action for a production mail server. Instead, you want the server to download the updates and notify you via the updates icon in the system tray. You can then investigate the updates and schedule appropriate downtime to apply them manually.

Finding Answers

This topic deserves special attention. One of our jobs is working in Tier 3 support for a large organization. The thing we respect the most about the administrators who actually run the system and handle the trouble tickets is that they do their homework prior to coming to us with a problem.

Too often techies make up an answer when they are not sure about something. Don't do that! When you are asked a question that you don't know the answer to, it is okay to say you don't know the answer—but make sure to follow that up by indicating that you will find the answer. Knowing the right resources (where to get answers) is just as important as the technical knowledge it takes to implement the answer. Key players in your organization will respect you much more when they know that you are willing to accept the limitations of your knowledge and have the appropriate resources to find the resolution to a problem or the answer to a question.

HELPFUL RESOURCES

Exchange Server has to be one of the most documented and discussed products (short of maybe Windows) that Microsoft produces. This means that most of the questions that we have about

Exchange Server can usually be answered via the right search or by looking in the right place. The most obvious place to start when you have a problem or a question is to perform an Internet search, but many other resources are available:

Exchange Server Documentation There is a world of free information on the Internet, but let's start right on the local hard disk of your Exchange Server or any place you have installed the admin tools. Microsoft has done an excellent job of providing better and better documentation for Exchange Server over the past few years. The Exchange Server 2016 documentation is comprehensive and so readable you will wonder if it is really from Microsoft. A link to the documentation can be found in the installation directory of Exchange Server. Look for the following file:

`C:\Program Files\Microsoft\Exchange Server\v15\Bin\ExchHelp.url`

You can also run it from the Microsoft Exchange Server 2016 folder on the Start menu. Either option will open a web browser that navigates to the TechNet reference library for Exchange Server.

Exchange Server Release Notes Another good resource for "I wish I had known that" types of things is the release notes. You should be able to find a link to the release notes here:

`C:\Program Files\Microsoft\Exchange Server\v15\`

Exchange Server Forums If you have a question for which you have done your due diligence in searching and researching the problem but you don't have an answer, it is time to ask the world. A good place to start is the Microsoft forums, also known as social.technet.microsoft.com. You can find the Exchange Server section here:

<http://social.technet.microsoft.com/forums/en-US/category/exchangeserver/>

When you post your question, please take a moment to think about what information the other readers are going to need to answer your question. Although you can post a vague question such as "Exchange is giving me an error," doing so is only going to result in (at best) delays while other forum participants have to request specific information from you. Instead, post the exact error message and any error codes you are seeing. Also, indicate, at minimum, what version of the software you are using (including service pack), the role of the server, and what operating system you are using.

You Had Me at EHLO This is the Microsoft Exchange Team's blog. This is the best site on the Internet for getting the inside scoop on how Exchange Server works, best practices, and the future of Exchange Server. You can read articles written by Exchange Server developers and Customer Support Services engineers. When changes to the product are announced, or customers request changes in the product, you will hear *first* from the product group engineers about the way they have chosen to deal with the issue.

<http://blogs.technet.com/b/exchange/>

MSExchange.Org Website One of the best sites on the Internet for free, easy-to-access content about Exchange Server is www.msexchange.org. The articles are written by Exchange Server gurus from all over the world and are usually in the form of easy-to-read and easy-to-follow tutorials. There is also a forums section where you can post questions or read other people's questions.

CALLING FOR SUPPORT

If your system is down or your operations are seriously hindered and you don't have a clue what to do next, it is time to call in the big guns. Sure, you should do some Internet searches to try to resolve your problem, but Internet newsgroups and forums are not the place to get support for business-critical issues.

Microsoft Product Support Services (PSS) is Microsoft's technical support organization. Its home page is <http://support.microsoft.com>. Professional support options (ranging from peer-to-peer support to telephone support) can be found at the following URL, where a web browser-based wizard guides you through your support options:

<https://gettechsupport.microsoft.com/default.aspx?locale=en-us&supportregion=en-us&pesid=14886>

If you do not have a Microsoft Premier agreement, Microsoft telephone support may seem to be a bit expensive, but believe me, when an Exchange server is down and the users are burning you in effigy in the company parking lot, a few hundred dollars for business hours support is cheap.

When you call and get a support technician on the phone, don't be surprised or offended if they start at the beginning and ask you a lot of elementary questions. They have to double-check everything you have done before they can look into more advanced problems. Frequently, one of these basic questions will help you locate a problem that you were convinced was more complicated than it really was. Though the beginning of the call may be underwhelming, the technician will stay with you on the phone until the problem is resolved or some kind of an acceptable resolution is put in place.

We always encourage people to call PSS if they truly need assistance. But PSS engineers are not mind readers, nor do they know every bit of Exchange Server code. You will do both yourself and the PSS engineer a big favor if you have all of your ducks in a row before you call. Do the following before you call:

- ◆ Attempt a graceful shutdown and restart of the server in question, if applicable.
- ◆ Perform a complete backup if possible.
- ◆ Have a complete, documented history of everything you have done to solve the problem. At the first sign of trouble, you should start keeping a chronological log of the things you did to fix the problem.
- ◆ Find out if you are allowed to initiate support sessions with remote support personnel through a tool like Skype for Business 2015 or WebEx.
- ◆ Be at a telephone that is physically at the server's console, or be in a place where you can access the server remotely via the Remote Desktop client. Your support call will be very brief if you cannot immediately begin checking things for the PSS engineer.
- ◆ Have the usernames and passwords that will provide you with the right level of administrative access. If you don't have those, have someone nearby who can log you in.
- ◆ Save copies of the event logs. Be prepared to send these to PSS if requested.

- ◆ Know the location of your most recent backup and how to access it when needed.
- ◆ Keep copies of all error messages. Don't paraphrase the message. Screen captures work great in this case. Pressing Alt+Print Scrn (or using the Snipping tool) and saving the screen capture as a file works great, too. We usually create a document with screen captures along with notes of what we were doing when we saw each message.

Be patient; telephone support is a terribly difficult job. A little kindness, patience, and understanding on your part will most certainly be returned by the PSS engineer.

Tools You Should Know

Out of the box, Exchange Server is an excellent product, but sometimes the base software that you install can use some assistance. Some of these tools are actually installed with Exchange Server, whereas you may need to download other tools.

PowerShell and the Exchange Management Shell Even here in the very first chapters, we are extolling the virtues of PowerShell. PowerShell enables some basic Windows management functions, such as managing event logs and services, to be performed via a command-line interface. This interface is simple to use and easy to learn, even for a GUI guy. The Exchange Server team pioneered the adoption of PowerShell when they built the entire Exchange Server 2007 management interface, known as the Exchange Management Shell (EMS), as an extension to PowerShell. Exchange Server 2013 and Exchange Server 2016 continue to follow this pattern.

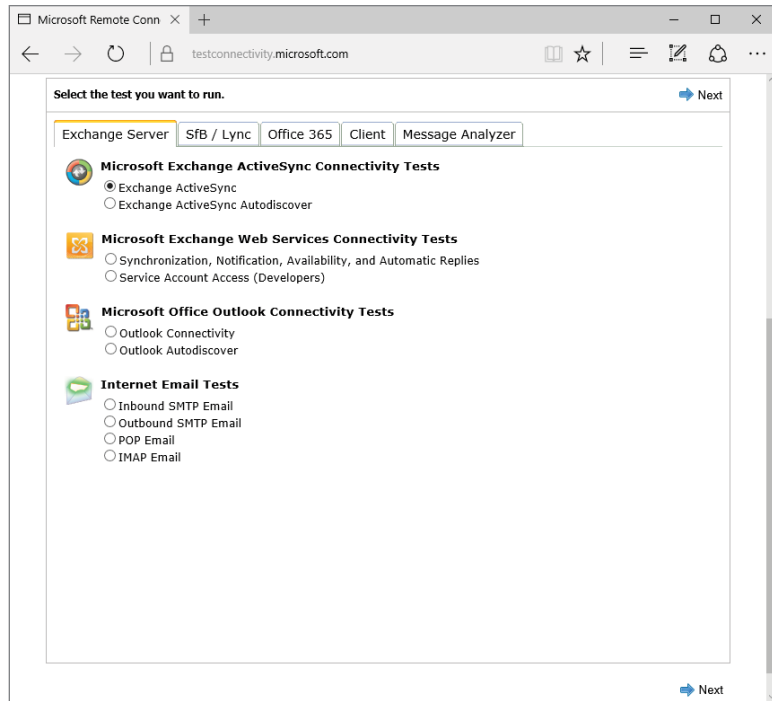
Although almost every chapter in this book will include at least some information about using EMS to perform Exchange Server management tasks, we have dedicated all of Chapter 5, “Introduction to PowerShell and the Exchange Management Shell,” to helping you learn your way around EMS.

Exchange Management Shell Test Cmdlets The Exchange Management Shell has a series of command-line tools that are very good for testing and diagnosing problems. These include tools for testing Outlook on the web connectivity, Unified Messaging connectivity, Outlook connectivity, and even mail flow. They are installed when you install the Exchange Server 2016 Management Tools. For more information, at the EMS prompt, enter **Get-Excommand test***.

Microsoft Remote Connectivity Analyzer (Previously Exchange Remote Connectivity Analyzer) Available at www.testexchangeconnectivity.com, the Remote Connectivity Analyzer is likely going to be the most useful tool in your troubleshooting arsenal. Initially started as a side project by two Microsoft employees, this website acts as the ultimate connectivity troubleshooting catch-all. The basic troubleshooting scenarios for Exchange Server 2016 (on-premises) are shown in Figure 1.7.

Those of you who have used “analyzers” from Microsoft in the past may remember the Exchange Best Practices Analyzer (ExBPA). The Remote Connectivity Analyzer should not be confused with the ExBPA. In fact, a new version of the ExBPA has not been released for Exchange Server since Exchange Server 2010.

FIGURE 1.7
Viewing the
Microsoft Remote
Connectivity
Analyzer



The Bottom Line

Understand email fundamentals. To gain the best advantage from Exchange Server 2016, you should have a good grounding in general email applications and principles.

Master It What two application models have email programs traditionally used? Which one does Exchange Server use? Can you name an example of the other model?

Identify email-administration duties. Installing an Exchange Server system is just the first part of the job. Once it's in place, it needs to be maintained. Be familiar with the various duties and concerns that will be involved with the care and feeding of Exchange Server.

Master It What are the various types of duties that a typical Exchange Server administrator will expect to perform?



Chapter 2

Introducing the Changes in Exchange Server 2016

Email clients used to be fairly simple and text based. Email servers had few connectivity options, no high-availability features, and no integrated directory. Then, beginning in the mid-1990s, we saw a big push toward providing email service to most of our user communities. We also saw email go from an occasionally used convenience to a business-critical tool. Business management and users demanded more features, better availability, and more connectivity options as the email client and server evolved.

Microsoft released Exchange Server 4.0 (the first version of Exchange Server) in 1996, and the product has been evolving ever since. Exchange Server 2016 is the eighth major release of the Exchange Server family and represents continued evolution of the product. The features and functions of this new release include not only features requested from many thousands of Microsoft's customers but also requirements shared internally at Microsoft by Microsoft Consulting Services and their own IT department, which supports more than 100,000 mailboxes.

We'll explore how some product features have evolved to this latest release, providing context for functionalities that were added, removed, modified, renamed, or reinvented. As of this writing, most Exchange Server customers are still using Exchange Server 2013 rather than Exchange Server 2016. Therefore, we'll focus on the changes that have been made to Exchange Server since Exchange Server 2013.

IN THIS CHAPTER, YOU WILL LEARN TO:

- ◆ Understand the changes in Exchange server architecture
- ◆ Understand the changes in the Exchange Server roles

Getting to Know Exchange Server 2016

It seems that we approach any new release of Exchange Server with a sense of both excitement and trepidation. We look forward to the new features and capabilities that are introduced with a newer version of the product. Certainly, the new site-resiliency features, compliance functionalities, resource management, management features, and security features will allow us to deliver better, more reliable messaging services to our end users.

On the other side of the coin is the feeling that we have to learn a whole new series of features inside and out so that we can better use them. Sure, we know Exchange Server 2013 pretty well, but there will be new details to learn with Exchange Server 2016. Sometimes we have to learn these implementation or management details the hard way.

However, this milestone in the evolution of Exchange Server is a good one. We can't help but be excited about learning about this new version and sharing what we have learned. We hope that you will feel the same sense of excitement. We have picked a top-ten list of new features that we like and hope that you will investigate further as you start to learn Exchange Server 2016. Some of these are summarized in this chapter and many of these you will find in more detail in later chapters.

- ◆ Simplicity of server roles: Mailbox and Edge Transport
- ◆ Proxy traffic from and to Exchange Server 2016
- ◆ Outlook on the web (formerly Outlook Web App)
- ◆ MAPI over HTTP as the default protocol
- ◆ Document collaboration with SharePoint 2016 and OneDrive for Business
- ◆ Wizard for hybrid Office 365 environments
- ◆ New conditions and actions for Data Loss Prevention (DLP) policies
- ◆ Public folder support for In-Place eDiscovery and In-Place Holds
- ◆ Compliance Search with eDiscovery
- ◆ Redesigned architecture for mailbox searches

LEARN THE EXCHANGE MANAGEMENT SHELL (AND WEAR SUNSCREEN!)

To those of you who have been around the Internet long enough to remember the “Wear Sunscreen” email, that was supposedly the 1997 commencement address to MIT given by Kurt Vonnegut but was in reality a column written by the *Chicago Tribune*'s Mary Schmich, we give you “Learn the Management Shell (and Wear Sunscreen)” to help you prepare for Exchange Server 2016, project management best practices, and the world in general:

- ◆ If we could offer you one important tip when learning Exchange Server 2016, it would be that you should get to know the Exchange Management Shell (EMS). Sure, it looks intimidating and nearly everything you will ever need to do is in the Exchange Admin Center. Many Exchange Server gurus will back us up on the value and usefulness of the EMS, whereas they might not agree with us on things such as using real-time block lists, making full backups daily, and keeping lots of free disk space available.
- ◆ Make regular Exchange Server data backups.
- ◆ Document.
- ◆ Don't believe everything you read from vendors; their job is to sell you things.

- ◆ Don't put off maintenance that might affect your uptime.
- ◆ If you get in trouble, call for help sooner rather than later. A few hundred dollars for a phone call to your vendor or Microsoft Product Support Services is better than a few days of downtime.
- ◆ Share your knowledge and configuration information with coworkers.
- ◆ Accept certain inalienable truths: disks will fail, servers will crash, users will complain, viruses will spread, and important messages will sometimes get caught in the spam filter.
- ◆ Get to know your users and communicate with them.
- ◆ Implement site resiliency and high availability for mailboxes *and* for public folder mailboxes.
- ◆ Make regular backups of your Active Directory.
- ◆ If a consultant is telling you something that you know in your gut is wrong, double-check their work or run their recommendation by another colleague. Second opinions and another set of eyes are almost always helpful.
- ◆ Think twice. Click once.

But trust me on the EMS.

In this chapter, we will cover the features of Exchange Server 2016 not only to give experienced Exchange Server administrators the proper perspective on Exchange Server 2016 but also to educate newly minted Exchange Server administrators on just how powerful Exchange Server has become. Some features we'll discuss in this chapter aren't brand new, but they are so key to the product and have been so greatly improved in this release that we are compelled to mention them at the outset.

Exchange Server Architecture

Over the last several releases, a number of significant changes have been made to the architecture of Exchange Server. These changes positively improve the performance and scalability of Exchange Server, but they also result in some pretty significant differences in the platform on which you support Exchange Server.

Windows Server 2012 R2 and Exchange Server 2016

Because of some of the underlying requirements of Exchange Server 2016, you must run Windows Server 2012 or Windows Server 2012 R2. The following editions of Windows Server will support Exchange Server 2016:

- ◆ Windows Server 2012 Standard Edition
- ◆ Windows Server 2012 Datacenter Edition

- ◆ Windows Server 2012 R2 Standard Edition
- ◆ Windows Server 2012 R2 Datacenter Edition

It may also be safe to assume that Exchange Server 2016 will also be supported on Windows Server 2016. However, at the time of this writing, Windows Server 2016 is still only available as a technical preview. Because of this, Exchange Server 2016 has not yet been qualified on Windows Server 2016.

Exchange Server 2016 also has several other requirements. These requirements include:

- ◆ Windows Management Framework 4.0
- ◆ Microsoft .NET Framework 4.5.2
- ◆ A forest function level of Windows Server 2008 or higher
- ◆ All domain controllers must be running Windows Server 2008 or later

The supported Outlook clients for Exchange 2016 include:

- ◆ Outlook 2016 with the latest service packs and updates
- ◆ Outlook 2013 with the latest service packs and updates
- ◆ Outlook 2010 with the latest service packs and updates
- ◆ Outlook for Mac for Office 365

The management tools for Exchange Server 2016 can be installed on a computer that has one of the following operating systems:

- ◆ Windows Server 2012 Standard or Datacenter
- ◆ Windows Server 2012 R2 Standard or Datacenter
- ◆ Windows 10 64-bit
- ◆ Windows 8.1 64-bit

Note that Exchange Server 2016 and Exchange Server 2007 cannot coexist in the same environment.

To install Exchange Server 2016 with Exchange Server 2010, the Exchange Server 2010 server must be running Update Rollup 11 for Exchange 2010 SP3 or later.

To install Exchange Server 2016 with Exchange Server 2013, Exchange 2013 Cumulative Update 10 or later must be installed on all Exchange Server 2013 servers in the organization.

Server Roles

Exchange Server 2013 had three server roles: the Client Access server role, the Edge Transport server role, and the Mailbox server role. In Exchange Server 2016, there are now just two server roles. The Client Access server role has been retired. Now, the two server roles are the Mailbox server role and the Edge Transport server role. The Mailbox server role includes all of the

components that a Client Access server role provided with Exchange Server 2013. The Mailbox server role now provides these services:

- ◆ Client Access protocols
- ◆ Transport service
- ◆ Mailbox databases
- ◆ Unified messaging

The Edge Transport server role is designed to enable you to deploy a messaging server in a perimeter network, outside of an Active Directory Domain Services (AD DS) environment. This assists in minimizing the attack surface of your Exchange environment. It also assists by adding a point of security for messages that include viruses and spam, keeping them out of the internal network.

Exchange Server 2016 also gives you the ability to proxy traffic from an Exchange Server 2013 environment, as well as from Exchange Server 2016 to Exchange Server 2013. This flexibility enables you to control the process of migrating to Exchange Server 2016, such as with a phased mailbox approach. It is also beneficial for interoperability between Exchange Server 2013 and Exchange Server 2016 because any mailbox server can proxy clients to the correct server, regardless of whether the server is running Exchange Server 2013 or Exchange Server 2016. We talk more about migrations and interoperability in Chapter 11, “Upgrades and Migrations to Exchange Server 2016 or Office 365.”

HIGH-AVAILABILITY DECISIONS

High-availability decisions do not need to be made at installation time. High availability for Exchange Server 2016 databases is added incrementally *after* the initial deployment of the Mailbox server. There is no clustered Mailbox server installation option; however, administrators create Database Availability Groups (DAGs) to implement high availability. High availability is discussed in detail in Chapter 20, “Creating and Managing Database Availability Groups.” Mailbox databases can be added to database availability groups at any point in the game. The databases can be removed from database availability groups as well, as needed. Essentially, the high-availability decisions can be done incrementally after a deployment has occurred and reversed if they no longer serve the needs of the organization. It is important to note that a DAG can contain only servers that run the same version of Exchange Server. Adding an Exchange Server 2016 to a DAG that contains Exchange Server 2013 servers is not supported, and vice versa.

THE MAILBOX SERVER ROLE

The Mailbox server role is responsible for so much, yet changes in the architecture have ensured that it requires few resources to perform all its necessary tasks. We will discuss, in later chapters, the database benefits with regard to the database schema and memory utilization in Exchange Server 2016. Recent improvements are designed to enhance the ability of a Mailbox server to do so much more with so much less.

Another very significant change in the Mailbox server role is the number of Client Access features that are now handled by this role. In Exchange Server 2016, a Mailbox server handles the data rendering for client requests, runs all of the client access protocols, and still maintains all mailboxes.

The Mailbox server role is responsible for the following functionality (this list isn't exhaustive):

- ◆ Hosts mailbox databases
- ◆ Hosts public folder database
- ◆ Provides transport-related services, including proxying (note that transport was originally handled by a Hub Transport server role that went away in Exchange Server 2013)
- ◆ Provides client connectivity for all clients (note that client access was handled by the Client Access server role in Exchange Server 2013 but is now handled by the Mailbox server role in Exchange Server 2016)

THE EDGE TRANSPORT SERVER ROLE

The amount of spam, malicious email, and viruses that some organizations receive is staggering. Even small organizations are receiving tens of thousands of pieces of spam, dozens of viruses, and hundreds of thousands of dictionary spamming attacks each week. Some organizations estimate that more than 90 percent of all inbound email is spam or other unwanted content. Keeping this unwanted content away from your Exchange servers is important. A common practice for messaging administrators is to employ additional layers of message hygiene and security. The first layer is usually some type of appliance or third-party SMTP software package that is installed in the organization's perimeter network. The problem with these third-party utilities is that the administrator has to become an expert on an additional technology. An easier method that some organizations choose is to use a cloud-based solution. The Exchange Online Protection (EOP) service from Microsoft is a popular cloud-based message protection solution.

Exchange Server 2016 includes a server role named Edge Transport. The role remains similar to the role from Exchange Server 2010 and Exchange Server 2013. The Edge Transport server role is recommended for perimeter networks outside of an AD DS environment. Although it is possible to install the Edge Transport role on a domain server, none of the Exchange services used for Edge Transport require AD DS. The Edge Transport server role uses Active Directory Lightweight Directory Services (AD LDS) to store configuration and recipient information.

An Edge Transport server will handle all inbound and outbound messaging traffic for a Mailbox server. This includes mail relay and smart host services for the Exchange environment. You can deploy multiple Edge Transport servers to enable redundancy and failover capabilities in the perimeter network. You can also load balance incoming messages by distributing the SMTP traffic to multiple Edge Transport servers.

IS THE EDGE TRANSPORT SERVER ROLE REQUIRED?

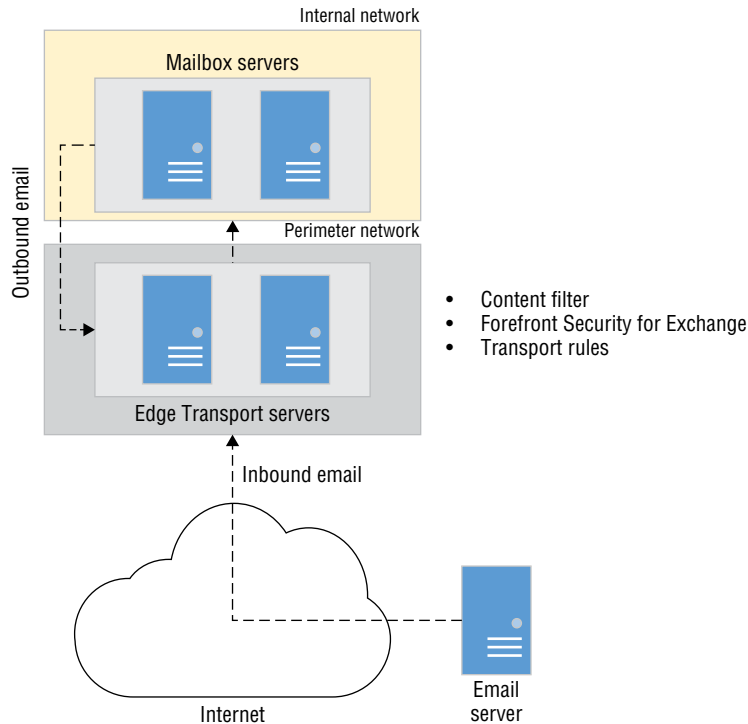
A common misconception is that the Edge Transport role is required for an Exchange Server organization. This is not the case, especially for organizations that choose to use a cloud-based message protection solution. Inbound email can be sent directly to the Mailbox server, or you can continue to use your existing third-party antispam/message-hygiene system to act as an inbound message relay for Exchange Server.

The Edge Transport server is a stand-alone message transport server that is managed using the EMS and the same basic management console that is used to manage Exchange Server 2016. A server functioning in an Edge Transport role should not be a member of the organization's internal Active Directory domain, although it can be part of a separate management forest used in a perimeter network.

Content filtering and Microsoft Forefront Security for Exchange are implemented on the Edge Transport server through content filtering and other antispam features. You can also run the features on the mailbox server if you do not have Edge Transport servers.

An example of how an organization might deploy an Edge Transport server is shown in Figure 2.1. Inbound email is first delivered to the Edge Transport servers that are located in the organization's perimeter network, where the message is inspected by the content filter, Forefront Security for Exchange, and any message transport rules. The inbound message is then sent on to the internal servers. Additionally, the Exchange Server Mailbox servers are configured to deliver mail, leaving the organization to the Edge Transport servers rather than configuring the internal servers to deliver mail directly to the Internet.

FIGURE 2.1
Deploying an Edge
Transport server



The Edge Transport server is a fully functional SMTP message-hygiene system with many of the same features that are found in expensive message-hygiene software packages and appliances. The following features are included:

- ◆ Per-user safe-sender, safe-recipient, and blocked-sender lists are automatically replicated from the user's mailbox to the Edge Transport server. Recipient filtering is enabled when valid recipients are synchronized to the Edge Transport server's local Active Directory Lightweight Directory Services (AD LDS) database.

- ◆ Sender and recipient filtering can be configured via administrator-controlled lists.
- ◆ Integrated Microsoft content filter is included for spam detection. Spam can be rejected, deleted, quarantined, or delivered to the user's Junk email.
- ◆ Multiple message-quarantines allow messages that are highly likely to be spam to be quarantined and sent to a quarantine mailbox on your Exchange server. A separate quarantine exists in the form of the user's Junk email folder for messages that are still tagged as spam but with a lower Spam Confidence Level.
- ◆ Microsoft Forefront Security for Exchange Server is available for the Edge Transport server when Enterprise client access licenses are used. However, this will be a short-lived solution, since Microsoft has announced that the entire suite of Forefront products is being decommissioned. Instead, many organizations use EOP or another third-party solution.
- ◆ Daily content filter and virus signature updates are available for organizations using Microsoft Forefront Security for Exchange Server.
- ◆ Real-time block lists and the IP Reputation Service allow an IP address to be checked to see if it is a known source of spam. Reputation filters can be updated on a daily basis.
- ◆ Sender ID filters allow for the verification of the mail server that sent a message and whether it is allowed to send mail for the message sender.
- ◆ Sender reputation filters allow a sender to be temporarily placed on a block list based on characteristics of mail coming from that sender, such as message content, sender ID verification, and sender behavior.

Client Connectivity

With Exchange Server 2013, Outlook clients connected to the Exchange Server by using RPC over HTTP (Outlook Anywhere). This enabled Outlook to connect to an Exchange server, regardless of its location, by using the Outlook Anywhere service.

Beginning with Exchange Server 2016, Outlook clients connect to the Exchange server by using MAPI over HTTP. RPC over HTTP is still available, but is official de-emphasized (meaning that it may not be included in future releases of Exchange Server). MAPI over HTTP is the default communication method between the client and the server.

MAPI over HTTP increases reliability and stability of the client connection. This protocol enables a higher level of visibility to errors that might occur between the client and server, as well as enhanced recoverability. MAPI over HTTP also includes support for a pause and resume function, which enables the clients to change networks while maintaining a connection to the Exchange Server. MAPI over HTTP can also reduce the total number of client connections, which can be helpful from a performance perspective.

While MAPI over HTTP is the default connection protocol for new Exchange Server 2016 environments, if you install Exchange Server 2016 in an environment with Exchange Server 2013, the protocol will not be used automatically. This is because MAPI over HTTP is not enabled by default in Exchange Server 2013 and was introduced with Exchange Server 2013 Service Pack 1.

Hybrid Improvements

Exchange Server 2016 can be implemented with Office 365 for a hybrid on-premises and cloud-based service. When configuring a hybrid organization with Exchange 2016, you will be

prompted to download the Hybrid Configuration Wizard. This wizard is included to assist configuring the hybrid environment.

The wizard has been updated for Exchange Server 2016 to include the following features:

- ◆ Easy updates for changes in Office 365 services
- ◆ Assists in troubleshooting a hybrid environment configuration
- ◆ Improved diagnostic information to resolve problems
- ◆ Support for both Exchange Server 2013 and 2016 hybrid environments

Hybrid deployments should be performed by using Azure Active Directory Connect (AAD Connect). AAD Connect provides functionality to synchronize multiple on-premises AD DS forests with a single Office 365 account.

In a hybrid environment, Exchange ActiveSync clients will be automatically directed to Office 365 if the user's mailbox is moved to the cloud. To support this automatic redirection, the ActiveSync client must support HTTP 451 redirects. After the client has been redirected, the Exchange profile on the device will be updated to use the new URL of the Exchange Online service. At this point, the client will not contact the on-premises environment for mailbox information.

OneDrive for Business Integration

With Exchange Server 2016 and SharePoint 2016, Outlook on the web users can link to and share documents that are stored in OneDrive for Business or on an on-premises SharePoint server. Instead of attaching a file to an email message, users can link to documents directly from Outlook on the web. Users can collaborate in an on-premises deployment just as they can with Office 365.

If a user receives a Word, Excel, or PowerPoint file that is stored in OneDrive for Business or SharePoint 2016, the recipient can view and edit the file directly from Outlook on the web. For an on-premises environment, a server must be running Office Online Server, which is in preview at the time of this writing, in the on-premises organization.

After editing the file within Outlook on the web, the recipient can save or upload the file to OneDrive.

Performance

The new architecture of Exchange Server 2016 combines the core features into a single server role. As part of that architecture, the search functionality has also been redesigned. In previous versions of Exchange Server, the searching functions were not fault-tolerant and were performed synchronously. In Exchange Server 2016, searching is performed asynchronously and is decentralized. Search functions are distributed across all Exchange Servers in the organization, and retries are attempted if servers are too busy.

The search scalability has also been improved. Previously, up to 5,000 mailboxes could be searched simultaneously from the web app. With Exchange Server 2016, this has increased to 10,000 mailboxes. When using the EMS, there is no limit to the number of mailboxes that can be searched.

Improved Policy and Compliance Features

Exchange Server 2016 has made significant improvements to both Data Loss Prevention (DLP) and eDiscovery.

DATA LOSS PREVENTION (DLP)

In Exchange Server 2016, transport rules have been updated with several new predicates and actions. Also, the coolest new feature to hit transport rules is DLP policies. DLP policies are designed to prevent users from sharing sensitive information with unauthorized users.

Every transport rule has three components: conditions, actions, and exceptions. The conditions specify under which circumstances the rule applies, whereas the exceptions specify under which conditions it will not apply. Exchange Server 2016 has the ability to identify, monitor, and protect 80 different types of sensitive information based on conditions and actions.

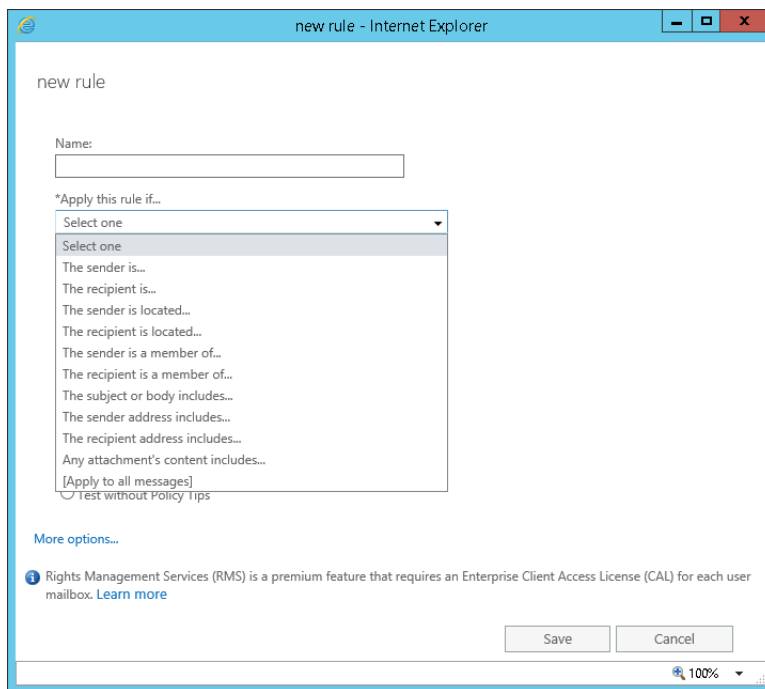
A new condition, “Any attachment has these properties, including any of these words,” will cause a trigger if an attached Office document contains the defined words. This condition enables you to integrate the transport rules with SharePoint, Windows Server 2012 R2 File Classification Infrastructure, or a third-party classification system.

A new action, “Notify the recipient with a message,” will send a customizable message to the recipient. For example, you can notify the recipient if the email was rejected or quarantined based on the contents.

The existing action “Generate incident report and send it to” has been updated so that the report can be messaged to multiple distribution lists.

The actions are the interesting part of the transport rule. Figure 2.2 shows the conditions on the New Rule window of the Transport Rule Wizard; this screen has three parts. The first part is checking on which object to take action, the second is simply checking the actions to take, and the third part specifies more details about the action.

FIGURE 2.2
Examining a
transport rule



eDISCOVERY AND PUBLIC FOLDERS

The market for third-party tools to support Exchange Server has grown rapidly. At one point, there were more than 60 third parties providing email archive solutions for Exchange Server. The sheer volume of email that users receive and their demand to keep historical email have made these tools very attractive.

Exchange Server 2010 introduced, and Exchange Server 2016 continues, a premium feature that allows for the integration of email archiving. The email archiving feature is actually a series of features that interact directly with the user's mailbox:

Archive Mailbox An archive mailbox is a secondary mailbox for a user that is used to store long-term email (archive email). An archive mailbox can be used in place of .pst files. Users can copy email messages from their primary mailbox to their archive mailbox. Archive mailboxes help users deal with large volumes of email while staying within mailbox size limits. The archive mailbox is defined on a user-by-user basis because not all users need an archive mailbox. The content in the archive mailbox can be accessed by users using the Outlook 2010 or later client or Outlook on the web.

Retention Policies Retention policies define the types of mail and how long the mail can be retained within the user's primary mailbox. Retention policies take the place of messaging records management (MRM) in Exchange Server 2007 and Exchange Server 2010. Retention policies can be defined to control when items are permanently deleted or when they are moved into the archive mailbox. With Outlook 2010 or later, end users can participate in the retention process by applying retention tags to messages or an entire folder.

eDiscovery (aka Multi-Mailbox and Federated Search) The eDiscovery features enables an authorized user to search for content across multiple data sources (both the user's "active" mailbox as well as their "personal archive mailbox") within an organization. You are able to search for information across Exchange, SharePoint, and Skype for Business archives, as well as use the eDiscovery Center in SharePoint 2013 to search for content in Exchange Server. Discovery managers can also export mailbox content to a .pst file from the SharePoint 2013 eDiscovery console. You can opt to use the Exchange Admin Center (EAC) to perform eDiscovery or opt to use SharePoint's eDiscovery Center. The eDiscovery Center offers some expanded capabilities, such as the ability to search and preserve content across multiple sources from a single console.

Exchange Server 2016 also introduces support for integrating public folders into eDiscovery. With In-Place eDiscovery, you can query public folders in the organization and put holds on public folders. Similar to placing a mailbox on hold, public folders support query-based and time-based holds. As of this writing, you can only search and hold all public folders. The ability to choose individual public folders to search and hold is expected in a later release.

In-Place Hold In-Place Hold enables an administrator to place a hold on a user's mailbox so that deleted and edited items are held during the hold period. This would be necessary in the event of legal action or an investigation regarding the conduct of one or more of your users.

Ultimately, the Exchange Server 2016 archiving and retention policies are intended to replace the messaging records-management features that were introduced in Exchange Server 2007.

eDiscovery and Compliance Search

A new feature of eDiscovery in Exchange Server 2016 is Compliance Search. Compliance Search is performed from the EMS, so there is no limit to the number of mailboxes that can be searched. For In-Place eDiscovery, you can search up to 10,000 mailboxes with a single search. Each Exchange Server organization can run up to two In-Place eDiscovery searches simultaneously.

To perform a Compliance Search, you must be assigned the Mailbox Search management role or be a member of the Discovery Management role group. The new EMS cmdlets available with Compliance Search are

- ◆ Get-ComplianceSearch
- ◆ New-ComplianceSearch
- ◆ Remove-ComplianceSearch
- ◆ Set-ComplianceSearch
- ◆ Start-ComplianceSearch
- ◆ Stop-ComplianceSearch

Message Transport Rules

Message transport rules are quite similar to Outlook rules and can even be created using a wizard similar to the one used to create Outlook rules. However, these rules are quite a bit more powerful and are run on Mailbox servers. Because all messages are processed by a Mailbox server regardless of whether they are inbound, outbound, or for local delivery, you can build powerful policies to control the messages and data that flow within your organization. Transport rules can also be defined at your organization's perimeter by using the Edge Transport server role in Exchange Server 2016.

New and Improved Outlook on the Web

Those of us who gushed when we saw the Outlook Web Access (OWA) interface in Exchange 2003 thought a web interface could not get much better. For Outlook on the web in Exchange 2013, the Exchange team started over from scratch to build a much more functional interface than ever before. For Exchange 2016, it has been updated and enhanced further! First, the name has changed! The new name is Outlook on the web. Here are some of the features in Outlook on the web:

- ◆ Platform-specific experiences for iOS and Android
- ◆ Premium Android experience with Chrome on Android version 4.2 or later
- ◆ Email improvements to the Inbox view and reading pane
- ◆ Contact linking with LinkedIn
- ◆ Updated calendar, including email reminders
- ◆ Search suggestions
- ◆ Thirteen new themes
- ◆ Preview URL links within messages

- ◆ Inline video playback from URLs
- ◆ Document collaboration with SharePoint 2016 and OneDrive for Business

Overview of Changes Since Exchange Server 2013

Since Exchange Server 2013, the primary changes to Exchange Server 2016 are

- ◆ Combined services (HT, CAS, MBX) in the Mailbox server role
- ◆ Integration with OneDrive and SharePoint 2016
- ◆ Additional policy and compliance features
- ◆ Outlook Web App redesigned as Outlook on the web

These are the key feature differences since Exchange Server 2013 and have been discussed in this chapter. Knowing some of the changes and introduction of features can be half of the battle to upgrading your knowledge on a newly released product.

Now, Where Did That Go?

As new and better functions and APIs have been introduced, naturally some functions are no longer emphasized or supported. We've already mentioned a few features that have been removed, but there are many more. There has been a lot of confusion surrounding what will continue to be supported in Exchange Server 2016 and what will no longer work. The phrase "no longer supported" itself tends to generate a lot of confusion because an unsupported function may continue to work because it has not truly been removed. Your mileage may vary when it comes to features that are no longer supported.

What's been removed from Exchange Server really depends on your perspective. Are you an Exchange Server 2010 expert? Is Exchange Server 2013 your comfort zone? We've broken down the next section of removed features based on your perspective.

Features No Longer Included

As Exchange Server has evolved into its current form, the code has experienced significant changes. Some features and APIs have been completely removed. Although most of these features will not affect the majority of Exchange Server deployments, you should keep them in mind and thoroughly evaluate your existing messaging environment to make sure you are not dependent on a feature that has no equivalent in Exchange Server 2016. If you require any of the features or APIs that were not carried over from Exchange Server 2010 or 2013, you may need to keep an older version of Exchange Server in operation.

EXCHANGE SERVER 2016 ESCHEWS EXCHANGE SERVER 2007

Only Exchange Server 2010 and Exchange Server 2013 can coexist with Exchange Server 2016 in the same organization. If you still require features provided by the Exchange Server 2007 platform, you will not be able to transition to Exchange Server 2016 until you can replace that particular feature requirement with newer software.

EXCHANGE SERVER 2010 FEATURES REMOVED FROM EXCHANGE SERVER 2016

The following features were included with Exchange Server 2010 but are no longer available in Exchange Server 2016:

- ◆ Unified Messaging directory lookups using Automatic Speech Recognition.
- ◆ Managed Folders for messaging retention management, including the Port Managed Folder Wizard.
- ◆ Antispam agents from the GUI. With Exchange Server 2016, antispam can be managed only from the EMS.
- ◆ Connection and Attachment filtering on Mailbox server roles. The only way to enable Connection Filtering is to use an Edge Transport server in a perimeter network.
- ◆ The ability to link a send-and-receive connector has been removed.
- ◆ Outlook Web App has been renamed to Outlook on the web. Additionally, spell check, customizable filters, message flags, chat contact lists, and search folders have been removed from the web client.
- ◆ Outlook 2003 and 2007 are not supported. Outlook clients must use either Outlook Anywhere (RPC over HTTP) or MAPI over HTTP.
- ◆ The Exchange Management Console and Exchange Control Panel have been replaced by the Exchange Admin Center.
- ◆ The Hub Transport and Unified Messaging server roles have been removed. Both server roles are included as features in the Mailbox server role.

EXCHANGE SERVER 2013 FEATURES REMOVED FROM EXCHANGE SERVER 2016

The following features are being de-emphasized with Exchange Server 2016 and may not be included in future versions:

- ◆ Third-party replication APIs.
- ◆ RPC over HTTP for client connections.
- ◆ Database Availability Group support for failover cluster administrative access points.
- ◆ Client Access server role. The functions of this role have been included in the Mailbox server role.
- ◆ The MAPI/CDO library has been replaced by Exchange Web Services, ActiveSync, and REST APIs.

Clearing Up Some Confusion

We mentioned earlier that Exchange has certainly been hyped a lot during the design and beta-testing process. This has generated a lot of buzz in the information technology industry, but this buzz has also generated a lot of confusion and some misinformation. Here we'll clear up the confusion by answering a few of the common questions about Exchange 2016.

Do I have to have two servers to run each of the server roles? In the days of Exchange Server 2010, many organizations deployed different roles to different servers in large organizations. Many administrators reserved the consolidated server approach for small environments. However, the performance capabilities of Exchange Server 2016 surpass the previous versions to such an extent that all services are run within the Mailbox server role.

Is there a 32-bit version of Exchange Server 2016? No, a 32-bit version of Exchange Server 2016 is not available.

Is the Edge Transport server required? No, Edge Transport servers are not required. You can use any third-party message-hygiene system in your perimeter network, you can direct inbound and outbound mail through your internal servers, or you can do both.

Is EMS knowledge required? Do I have to learn scripting? Most common administrative tasks can be performed through the Exchange Admin Center web-based interface. Command-line management and scripting for Exchange Server 2016 have been greatly improved through the use of the EMS. Many tasks are simpler or more powerful through the EMS, but it is not necessary to learn scripting in order to start working with Exchange Server 2016. We strongly encourage you to get to know many of the powerful features of the EMS as you get comfortable with Exchange Server 2016. A number of advanced administration tasks do not have a graphical user interface option.

What is happening with public folders? The use of public folders with Exchange Server 2016 is still available and supported. However, for years, there has been talk about moving away from public folders, potentially removing support for them at some point. At the time of this writing, there isn't any information to indicate that this is coming soon (or coming at all). But you may want to examine your public folder applications with an eye toward migrating them to systems such as Microsoft SharePoint Server 2016 to take advantage of the latest collaboration features. Also, remember that the traditional public folder databases are no longer available in Exchange Server 2016 and that you must now store all public folders in a public folder mailbox.

The Bottom Line

Understand the key changes in Exchange Server 2016. Significant updates were made to the Exchange Server 2016 architecture to continue the improvement to the scalability, security, and stability. The Mailbox role handles mailboxes, public folders, transport, and client connectivity. Compliance features, such as compliance search and eDiscovery, are greatly enhanced and simplified. The disk I/O requirements continue to be reduced, enabling organizations to run their Exchange servers on lower-performing storage.

Master It You are planning your email data storage strategy, especially for long-term storage. You want to minimize or eliminate the use of .pst files. Which technology should you use to maintain email data indefinitely?

Understand the Mailbox role's expanded duties. Over the last couple of versions of Exchange Server, the Exchange server roles have been updated. In each version, a server role was consolidated, enabling organizations to reduce their server footprint and simplify their environments.

Master It You are planning a training session for your junior administrators to prepare them in their SMTP connectivity troubleshooting tasks. Which server role should you recommend they inspect when attempting to troubleshoot email delivery problems?

Chapter 3

Understanding Availability, Recovery, and Compliance

The modern business world is getting more complex, not less; email in turn evolves to keep up. As an Exchange Server administrator or implementer, you need to know more about a wider variety of topics without losing your core competency in Exchange Server.

IN THIS CHAPTER, YOU WILL LEARN TO:

- ◆ Distinguish between availability, backup and recovery, and disaster recovery
- ◆ Determine the best option for disaster recovery
- ◆ Distinguish between the different types of availability meant by the term *high availability*
- ◆ Implement the four pillars of compliance and governance activities

Changing from a Technology to a Business Viewpoint

You've probably heard the old proverb that "every cloud has a silver lining." It can be a comfort to know that good can usually be found during even the worst occasions. When a mailbox database server's RAID controller goes bad and corrupts the drive array containing the executive mailboxes, you have the opportunity to validate your backup strategy and demonstrate that it works perfectly under pressure.

However, the unacknowledged corollary is Murphy's Law: "Anything that can go wrong will go wrong." Every feature, functionality, and component that is added to a messaging infrastructure increases complexity and the number of potential failures. If you think for a moment about the spread of email and how it has changed from a luxury to a utility, you can see that electronic messaging administrators have become victims of their own success.

Gone are the days where you simply had to worry about editing and publishing the correct DNS records for your domains, provisioning and configuring your T1 routers, and wrestling with server hardware. Today's challenges involve meeting more goals, supporting more complex environments, meeting business requirements, and analyzing risks. These are common scenarios:

- ◆ Ensuring that mailbox servers have the proper storage back-end design to allow backups to happen within a defined window
- ◆ Ensuring that your users continue to have access to their mailboxes even when a server fails, a flaky router takes a site offline, or power fails for an entire rack of servers

- ◆ Ensuring that a plan exists for quick recovery and restoration of your core messaging capabilities when the storage is offline or corrupt
- ◆ Ensuring that the messages users send to external clients are in compliance with all business policies and regulations
- ◆ Determining the risks associated with failing to provide disaster-recovery plans and the risks associated with a failure to meet service-level agreements
- ◆ Balancing business costs versus risks associated with providing recovery, ensuring compliance, and providing a specified level of service

What's in a Name?

Backup and recovery, high availability, disaster recovery, and compliance and governance—you have likely heard of these many times. Each plays a role in the overall protection strategy for your organization's data.

Each of these topics must be evaluated by every modern Exchange Server administrator and professional, along with appropriate business stakeholders, even if they are not actively addressed in every deployment of Exchange Server 2016. When you do need to address them in your planning, Exchange Server 2016 provides a variety of options to ensure that the deployment meets the particular needs of your business. One size and one set of capabilities do not fit all organizations. To make the best use of the tools that Exchange Server gives you, you must clearly understand the problems that each capability is designed to solve. It doesn't help to use a screwdriver when you need a hammer—and you can't solve a disaster-recovery problem by using an eDiscovery search.

In this section, a common vocabulary will be presented for discussing these topics. This will enable you to get the most from our discussions of the new features and functionality in Exchange Server 2016 that are covered in later chapters. You should clearly understand how Microsoft intended Exchange Server 2016's features to be deployed and used, so that you have confidence that they will meet your business goals.

Backup and Recovery

Let us begin with a topic that is one of the core tasks for any IT administrator, not just Exchange Server administrators: backup and recovery.

Backup is the process of preserving one or more point-in-time copies of a set of data, regardless of the number of copies, frequency and schedule, or media type used to store them.

As an administrator, you need to make sure your backups include all of the components you need to get Exchange Server services up and running again. That means more than just the databases. You should also consider the following components:

- ◆ **Active Directory Domain Services.** Exchange Server relies on Active Directory, so it is critical that Active Directory is highly available and backed up. Your Active Directory administrators probably handle this. But no matter who handles it, you should ensure that the backups are in place.
- ◆ **Operating system for the Exchange servers (System State as a minimum).** Prior to virtualization, backing up the Exchange server operating systems was quite important because building a new physical server (or rebuilding a physical server) from scratch was

time-consuming. Today, with virtualization, building a new server is quite fast. Some organizations opt to deploy new servers and forgo the backup of the operating system for some servers. However, without a backup of the operating system or system state, you will lose customizations such as in IIS and the Registry.

- ◆ **File system.** The file system has log files, configuration files, and other data that can be helpful in a disaster-recovery situation.
- ◆ **Database and database log files.** The Exchange databases are a critical piece of your backups because all of the email data is stored in the databases!

As you can see, backing up all of the components can quickly become complicated. It is important to have the right backup tools at your disposal. As part of your disaster-recovery planning, you should look at the available tools, including third-party tools, to figure out which tools best meet your requirements and provide the best administrative experience.

With Exchange Server, there are four main types of database backups:

Full Backups (Normal) Full backups capture an entire set of target data; in early versions of Exchange Server, this is a storage group with the transaction log files and all the associated mailbox databases and files. Beginning with Exchange Server 2010 and continuing in Exchange Server 2013 and Exchange Server 2016, each mailbox database is a separate backup target, since there is now an enforced 1:1 relationship between mailbox databases and transaction logs (it was “strongly recommended” in earlier versions). Full backups take the most time to perform and use the most space. If circular logging is disabled for a mailbox database, full backups must be executed on a regular basis. A successful full backup informs Exchange Server that the databases and transaction logs have been preserved and that saved transaction logs can be purged. Circular logging will be discussed in more depth later.

Copy Backups Copy backups are exactly like full backups, except that saved transaction logs are not purged.

Incremental Backups Incremental backups capture only a partial set of the target data—specifically, the data that has changed since either the last full backup or the last incremental backup. For Exchange Server, this means any new transaction logs. Incremental backups are designed to *minimize how often* full backups are performed, as well as *minimize the space used* by any particular backup set. As a result, a backup set that includes incremental backups can be more time-consuming and fragile to restore; successful recovery includes first recovering the latest full backup and then each successive incremental backup. Incremental backups also instruct Exchange Server to purge the saved transaction logs after the backup is complete. Incremental backups are not available when circular logging is enabled.

Differential Backups Differential backups also capture only a partial set of the target data—specifically, the data that has changed since the last full backup. No other backups (incremental or differential) are considered. For Exchange Server, this means any transaction logs generated since the last full backup. Differential backups are designed to *minimize how many* recovery operations you have to perform in order to fully restore a set of data. In turn, differential backups use more space than incremental backups, but they can be recovered more quickly and with fewer opportunities for data corruption; successful recovery includes first recovering the latest full backup and then the latest differential backup. A differential backup does not purge saved transaction logs. Differential backups are not available when circular logging is enabled.

Also known as restoration, recovery is the process of taking one or more sets of the data preserved through backups and making it once again accessible to administrators, applications, and/or end users. Most recovery jobs require the restoration of multiple sets of backup data, especially when incremental and differential backups are in use. Two metrics are used to determine if the recovery time and the amount of data recovered are acceptable:

Recovery Time Objective Recovery Time Objective (RTO) is a metric commonly used to help define successful backup and restore processes. The RTO defines the time window in which you must restore Exchange Server services and messaging data after an adverse event. You may have multiple tiers of data and service, in which case it could be appropriate to have a separate RTO for each tier. Often, the RTO is a component of (ideally, an input into, but that's not always the case) your service-level agreements. As a result, the RTO is a critical factor in the design of Exchange Server mailbox-database storage systems; it's a bad idea to design or provision mailbox databases that are larger than you can restore within your RTO.

Recovery Point Objective Recovery Point Objective (RPO) is a metric that goes hand in hand with the RTO. While the RTO measures a time frame, the RPO sets a benchmark for the maximum amount of data (typically measured in hours) you can afford to lose. Again, multiple tiers of service and data often have separate RPOs. The RPO helps drive the backup frequency and schedule. It's worth noting that this metric makes an explicit assumption that all data within a given category is equally valuable; that's obviously not true, which is why it is important to properly establish your categories. Remember, though, if you have too many classes or categories, you'll just have confusion.

One thing to note about Exchange Server 2016 databases is that they support only online backups and restores created through the Windows Volume Shadow Copy Service (VSS). VSS provides several advantages compared to other backup methods, including the ability to integrate with third-party storage systems to speed up the backup and recovery processes. The most important benefit VSS gives, though, is that it ensures that the Exchange Server information store flushes all pending writes consistently, ensuring that a backup dataset can be cleanly recovered.

We will use the phrase “backup set” several times. A *backup set* is a copy of all of the various backups that are required to perform a particular recovery. This will almost always include at least the last full backup and may include one or more incremental or differential backups

HOW MUCH DATA GETS COPIED?

One thing that Volume Shadow Copy Service does not natively provide is the ability to reduce the amount of data that must be copied during a backup operation. VSS simply creates either a permanent or temporary replica (depending on how the invoking application requested the replica be created) of the disk volume; it's then up to the application to sort out the appropriate files and folders that make up the dataset. Usually, this is the entire disk volume, but depending on the selected *VSS writers* it may only be a portion of a disk volume or specific files on a disk volume. Many Exchange Server-aware backup applications simply copy the various transaction log files and mailbox database files to the backup server.

Some applications, however, are a bit more intelligent; they keep track of which blocks have changed in the target files since the last backup interval. These applications can copy just those changed blocks to the backup dataset—typically some percentage of the blocks in the mailbox database file as well as all the new transaction log files—thus reducing the amount of data that needs to travel over the network and be stored. Block-level backups help strike a good balance between storage, speed, and reliability. As you go forward with VSS-aware Exchange Server-compatible backup solutions, be sure to investigate whether they offer this feature. Microsoft's System Center Data Protection Manager does offer this feature.

Disaster Recovery

Regular backups are important; the ability to successfully restore them is even more important. This capability is a key part of your extended arsenal for problem situations. Restoring the occasional backup is fairly straightforward but assumes that you have a functional Exchange server and the dependent network infrastructures. What do you do if an entire site or datacenter goes down and your recovery operations extend beyond a single Exchange Server mailbox database? The answer to this question is a broad topic that can fill many books, blog postings, and websites of its own.

Disaster recovery (DR) is the practice of ensuring that critical services can be restored when some disaster or event causes large-scale or long-term outage. A successful DR plan requires the identification of critical services, dependencies, and data, creation of documentation that lists the necessary tasks to re-create and restore them, and modification of the relevant policies and processes within your organization to support the DR plan.

It's not enough to consider how to rebuild Exchange servers and restore Exchange Server mailbox databases. Exchange Server is a complex application with many dependencies, so your plans need to accommodate the following issues:

Network Dependencies These include subnets, IP address assignments, DNS, load balancers, DHCP services, switch configurations, network/Internet access, and router configurations. Are you rebuilding your services to have the same IP addresses or new ones? Whatever you decide, you'll need to make sure that required services and clients can reach the Exchange servers.

Active Directory Services These include associated DNS zones and records. Exchange Server cannot function without reliable access to global catalog servers and other domain controllers. Which forests and domains hold objects Exchange Server will need to reference? Does your existing replication configuration meet those needs during a DR scenario? What would happen if an Active Directory user account that was associated with a mailbox was accidentally deleted?

Third-Party Applications These include monitoring, backup, archival, or other programs and services that require messaging services or interact with those services. Don't just blindly catalog everything in production; be sure these systems are also being addressed as part of the disaster-recovery plan.

There's a blurry line between disaster recovery and the associated concept of *business continuity* (also called *business continuance*). Business continuity (BC) is the ability of your organization to continue providing some minimum set of operations and services necessary to stay in business during a large-scale outage, such as during a regional event or natural disaster (for example, a hurricane or earthquake). In a business continuity plan, your organization will identify and prioritize the most critical services and capabilities that need to provide at least some level of operational capacity as soon as possible, even without full access to data or applications.

It's important to note that the business continuity plan is designed and implemented alongside your disaster-recovery efforts. In many organizations, they will be maintained by two separate groups of professionals. It is imperative that these groups should have good lines of communication in place.



Real World Scenario

DRAWING THE LINE BETWEEN DISASTER RECOVERY AND BUSINESS CONTINUITY

There's a lot of confusion over exactly how disaster recovery and business continuity relate to each other. We have good news and bad news: The good news is that it's a simple relationship. The bad news is, "It depends."

Both types of plans are ultimately aimed at the goal of repairing the damage caused by extended outages. The biggest difference is the scope; many business-continuity plans focus very little on technology and look instead at overall business processes. In contrast, disaster-recovery plans of necessity have to be concerned with the finer details of IT administration. The reality is that both levels of focus are often needed—and must be handled in parallel, with coordination, and in support of any additional ongoing crisis management.

We'll try to clarify the difference by providing an example. Acme Inc. is a national manufacturer and supplier of various goods, mainly to wholesale distributors but with a small and thriving mail-order retail department for the occasional customer who needs quality Acme products but has no convenient retail outlet in their locale. Acme's main call center has a small number of permanent staff but a large number of contract call center operators.

Unfortunately, Acme's main order fulfillment center—for both bulk wholesale orders, as well as the relatively small amount of mail-order traffic—gets hit by a large fragment in a meteor shower, causing a fire that rapidly transforms the entire site into smoking rubble even as all personnel are safely evacuated. The call center and supporting datacenter are completely destroyed and, conservatively, will take several months to fully rebuild. Obviously, Acme is going to suffer some sort of setback, but with proper planning they can minimize the effects. What types of actions would Acme's BC and DR plans each be taking?

Acme's Business Continuity Plan Acme is concerned with getting the minimum level of operational function back online as quickly as possible. In this case, it's going to take a while before they can resume call center operations. Their immediate needs are to establish at least some level of messaging support for the temporary call center workers the BC plan brings in. Their BC plan does not assume that they will have in-house capability, so it makes provisions—if required—to use hosted Exchange Server services as a short-term stopgap so that communications with customers and wholesalers will proceed until Acme's IT staff can bring up sufficient Exchange servers to switch back to on-premises services.