Microsoft

Microsoft Exchange Server 2013 Databases, Services, & Management

William R. Stanek
Author and Series Editor

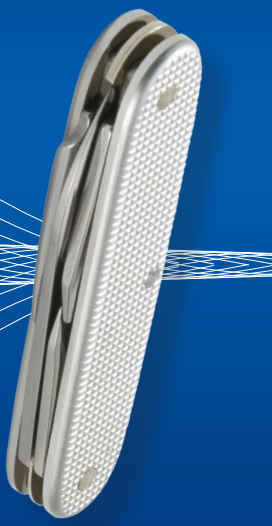

Pocket Consultant

Microsoft

Microsoft Exchange Server 2013: Databases, Services, & Management

Pocket Consultant

William Stanek
Author and Series Editor

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

Copyright © 2013 by William R. Stanek

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

Library of Congress Control Number: 2013949891
ISBN: 978-0-7356-8175-0

Printed and bound in the United States of America.

Second Printing: June 2014

Microsoft Press books are available through booksellers and distributors worldwide. If you need support related to this book, email Microsoft Press Book Support at mspinput@microsoft.com. Please tell us what you think of this book at http://www.microsoft.com/learning/booksurvey.

Microsoft and the trademarks listed at http://www.microsoft.com/en-us/legal/intellectualproperty/trademarks/en-us.aspx are trademarks of the Microsoft group of companies. All other marks are property of their respective owners.

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

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

Acquisitions Editor: Anne Hamilton
Developmental Editor: Karen Szall
Project Editor: Karen Szall
Editorial Production: Online Training Solutions, Inc. (OTSI)
Technical Reviewer: Todd Meister; Technical Review services provided by Content Master, a member of CM Group, Ltd.
Copyeditor: Denise Bankaitis (OTSI)
Indexer: Krista Wall (OTSI)
Cover: Best & Company Design

Acknowledgments

To my readers—*Microsoft Exchange Server 2013 Pocket Consultant: Databases, Services, & Management* is my 42nd book for Microsoft Press. Thank you for being there with me through many books and many years.

To my wife—for many years, through many books, many millions of words, and many thousands of pages she's been there, providing support and encouragement and making every place we've lived a home.

To my kids—for helping me see the world in new ways, for having exceptional patience and boundless love, and for making every day an adventure.

To Anne, Karen, Martin, Lucinda, Juliana, and many others who've helped out in ways both large and small.

Special thanks to my son Will for not only installing and managing my extensive dev lab for all my books since *Windows 8 Pocket Consultant* but for also performing check reads of all those books as well.

—William R. Stanek

Contents

Introduction *xiii*

Chapter 1 Microsoft Exchange organizations: the essentials 1

Understanding Exchange Server 2013 organizations 2
- Organizational architecture 2
- Front-end transport 4
- Back-end transport 6

Site-based and group-based routing 8
- Routing boundaries 8
- IP site links 9
- On-premises, online, and cross-premises routing 12

Understanding data storage in Exchange Server 2013 13
- Working with the Active Directory data store 14
- Working with the Exchange store 16

Chapter 2 Managing data and availability groups 25

Navigating the Information Store 25
- Basic database options 26
- High availability database options 28
- Working with Active Manager 30
- Understanding managed availability 32

Creating and managing database availability groups 34
- Pre-staging and preparing for database availability groups 34
- Creating database availability groups 39
- Managing availability group membership 42

What do you think of this book? We want to hear from you!

Microsoft is interested in hearing your feedback so we can continually improve our books and learning resources for you. To participate in a brief online survey, please visit:

microsoft.com/learning/booksurvey

Managing database availability group networks 45
Changing availability group network settings 49
Configuring database availability group properties 52
Removing servers from a database availability group 54
Removing database availability groups 54

Maintaining database availability groups 55
Switching over servers and databases 55
Checking continuous replication status 58
Restoring operations after a DAG member failure 59

Chapter 3 Exchange database administration 63

Working with active mailbox databases. 63
Understanding mailbox databases 64
Preparing for automatic reseed 65
Creating mailbox databases 66
Setting the default offline address book 70
Setting mailbox database limits and deletion retention 71
Recovering deleted mailboxes 75
Recovering deleted items from mailbox databases 77

Working with mailbox database copies 78
Creating mailbox database copies 79
Setting replay, truncation, and preference values for database copies 81
Suspending and resuming replication 83
Activating lagged database copies 84
Updating mailbox database copies 87
Monitoring database replication status 90
Removing database copies 94

Managing mailbox databases 95
Mounting and dismounting databases 95
Setting the maintenance interval 98
Moving databases 100
Renaming databases 102
Deleting databases 103

Content indexing 103
Understanding indexing 103
Managing Exchange Store Search 104
Troubleshooting indexing 106

Chapter 4 Configuring transport services 107

Working with SMTP connectors, sites, and links 108
Connecting source and destination servers 108
Managing Active Directory site details 109
Managing Active Directory site link details 112
Creating Send connectors 114
Viewing and managing Send connectors 122
Configuring Send connector DNS lookups 125
Setting Send connector limits 126
Creating Receive connectors 128
Viewing and managing Receive connectors 135
Creating Inbound and Outbound connectors with Exchange Online 139

Configuring transport limits 141
Setting organizational transport limits 142
Setting connector transport limits 143
Setting server transport limits 144

Completing Transport services setup 147
Configuring the postmaster address and mailbox 147
Configuring shadow redundancy 148
Configuring Safety Net 153
Enabling anti-spam features 154
Subscribing Edge Transport servers 156

Chapter 5 Managing and maintaining mail flow 163

Managing message pickup, replay, throttling, and back pressure 163
Understanding message pickup and replay 164
Configuring and moving the Pickup and Replay directories 165

Changing the message processing speed 166
Configuring messaging limits for the Pickup directory 167
Configuring message throttling 168
Understanding back pressure 169
Creating and managing accepted domains 170
Understanding accepted domains, authoritative domains, and relay domains 170
Viewing accepted domains 171
Creating accepted domains 173
Changing the accepted domain type and identifier 174
Removing accepted domains 176
Creating and managing email address policies 176
Viewing email address policies 176
Creating email address policies 178
Editing and applying email address policies 182
Removing email address policies 183
Configuring journal rules 184
Working with journal rules 184
Setting the NDR journaling mailbox 184
Creating journal rules 185
Managing journal rules 186
Creating and managing remote domains 186
Viewing remote domains 186
Creating remote domains 187
Configuring messaging options for remote domains 187
Removing remote domains 189
Configuring anti-spam and message filtering options 189
Filtering spam and other unwanted mail by sender 190
Filtering spam and other unwanted email by recipient 193
Filtering connections with IP block lists 195
Defining block list exceptions and global allow/block lists 201
Preventing internal servers from being filtered 205
Configuring transport rules 205
Understanding transport rules 205
Creating transport rules 206
Managing transport rules 208

Chapter 6 Managing client access 209

Mastering Outlook Web App essentials. 210
- Getting started with Outlook Web App 210
- Connecting to mailboxes and public folder data over the web 212
- Working with Outlook Web App 213
- Enabling and disabling web access for users 216
- Troubleshooting Outlook Web App 217

Managing web and mobile access 221
- Using Outlook Web App and Exchange ActiveSync with IIS 221
- Working with virtual directories and web applications 222
- Enabling and disabling Outlook Web App features 226
- Configuring ports, IP addresses, and host names used by websites 230
- Enabling SSL on websites 232
- Restricting incoming connections and setting time-out values 235
- Redirecting users to alternate URLs 236
- Controlling access to the HTTP server 237
- Throttling Client Access 241
- Starting, stopping, and restarting websites 244
- Configuring URLs and authentication for the OAB 244
- Configuring URLs and authentication for OWA 246
- Configuring URLs and authentication for Exchange ActiveSync 247
- Configuring URLs and authentication for ECP 248

Configuring POP3 and IMAP4 249
- Enabling the Exchange POP3 and IMAP4 services 250
- Configuring POP3 and IMAP4 bindings 252
- Configuring POP3 and IMAP4 authentication 253
- Configuring connection settings for POP3 and IMAP4 254
- Configuring message retrieval settings for POP3 and IMAP4 256

Managing Outlook Anywhere 257
- Working with Outlook Anywhere 257
- Configuring URLs and authentication for Outlook Anywhere 258

Chapter 7 Managing mobile messaging 261

Mastering mobile device and wireless access essentials. 261
Using Exchange ActiveSync and Outlook Web App for Devices 262
Managing Exchange ActiveSync and Outlook Web App for Devices 262
Moving from remote mail to Outlook Anywhere 263

Managing Exchange Server features for mobile devices 267
Using Autodiscover 268
Using Direct Push 271
Using remote device wipe 272
Using password recovery 276
Configuring direct file access 277
Configuring remote file access 282
Using WebReady Document Viewing 283

Working with mobile devices and device policies 285
Viewing existing mobile device mailbox policies 285
Creating mobile device mailbox policies 288
Optimizing mobile device mailbox policies 291
Assigning mobile device mailbox policies 293
Removing mobile device mailbox policies 294

Managing device access 295

Chapter 8 Exchange Server 2013 maintenance, monitoring, and queuing 299

Performing tracking and logging activities in an organization. 299
Using message tracking 299
Using protocol logging 307
Using connectivity logging 314

Monitoring events, services, servers, and resource usage 317
Viewing events 317
Managing essential services 320
Monitoring Exchange messaging components 321
Using performance alerting 323

Working with queues 328
Understanding Exchange queues 328
Accessing the Queue Viewer 330
Managing queues 331
Understanding queue summaries and queue states 331
Refreshing the queue view 332
Working with messages in queues 333
Forcing connections to queues 334
Suspending and resuming queues 334
Deleting messages from queues 335

Chapter 9 Troubleshooting Exchange Server 2013 337

Troubleshooting essentials 337
Tracking server health 337
Tracking user and workload throttling 342
Tracking configuration changes 343
Testing service health, mail flow, replication, and more 344
Diagnosing and resolving problems 348
Identifying recovery actions 348
Identifying responders 350
Identifying monitors 352
Identifying probes 353
Viewing error messages for probes 354
Tracing probe errors 356
Using Log Parser Studio 359
Getting started with Log Parser Studio 359
Performing queries in Log Parser Studio 360

Index *363*

About the author *381*

What do you think of this book? We want to hear from you!

Microsoft is interested in hearing your feedback so we can continually improve our books and learning resources for you. To participate in a brief online survey, please visit:

microsoft.com/learning/booksurvey

Introduction

Microsoft Exchange Server 2013 Pocket Consultant: Databases, Services, & Management is designed to be a concise and compulsively usable resource for Exchange Server 2013 administrators. This is a resource guide that you'll want on your desk at all times. The book covers everything you need to perform the core administrative tasks for Exchange databases, transport services, mail flow, and Client Access servers, whether your servers are running on Windows Server 2012 or Windows Server 2008 R2. Because the focus of this book is on giving you maximum value in a pocket-size guide, you don't have to wade through hundreds of pages of extraneous information to find what you're looking for. Instead, you'll easily find exactly what you need to get the job done.

This book zeroes in on daily administrative procedures, frequently performed tasks, documented examples, and options that are representative although not necessarily inclusive. One of my goals is to keep the content so concise that the book remains compact and easy to navigate while at the same time ensuring that the book is packed with as much information as possible. Thus, instead of a hefty 1,000-page tome or a lightweight 100-page quick reference, you get a valuable resource guide that can help you quickly and easily perform common tasks, and solve problems.

Although you might not install Exchange Server 2013 on touch-enabled computers, you can still *manage* Exchange Server 2013 from your touch-enabled computers; therefore, understanding the touch UI in addition to the revised interfaces options will be crucial to your success. For this reason, I discuss both the touch UI and the traditional mouse and keyboard techniques throughout this book.

When you are working with touch–enabled computers, you can manipulate on-screen elements in ways that weren't possible previously. You can enter text by using the on-screen keyboard and also in the following ways:

- **Tap** Tap an item by touching it with your finger. A tap or double-tap of elements on the screen generally is the equivalent of a mouse click or double-click.
- **Press and hold** Press your finger down and leave it there for a few seconds. Pressing and holding elements on the screen generally is the equivalent of a right-click.
- **Swipe to select** Slide an item a short distance in the opposite direction compared to how the page scrolls. This selects the items and also might bring up related commands. If pressing and holding doesn't display commands and options for an item, try using swipe to select instead.
- **Swipe from edge (slide in from edge)** Starting from the edge of the screen, swipe or slide in. Sliding in from the right edge opens the Charms panel. Sliding in from the left edge shows open apps and allows you to easily switch between them. Sliding in from the top or bottom edge shows commands for the active element.

- **Pinch** Touch an item with two or more fingers and then move the fingers toward each other. Pinching zooms in or shows less information.
- **Stretch** Touch an item with two or more fingers and then move the fingers away from each other. Stretching zooms out or shows more information.

As you've probably noticed, a great deal of information about Exchange Server 2013 is available on the web and in other printed books. You can find tutorials, reference sites, discussion groups, and more to make using Exchange Server 2013 easier. However, the advantage of reading this book is that much of the information you need to learn about Exchange Server 2013 is organized in one place and presented in a straightforward and orderly fashion. This book has everything you need to master Exchange databases, transport services, mail flow, and Client Access servers.

In this book, I teach you how features work, why they work in the way that they do, and how to customize features to meet your needs. I also offer specific examples of how certain features can meet your needs, and how you can use other features to troubleshoot and resolve issues you might have. In addition, this book provides tips, best practices, and examples of how to optimize Exchange Server 2013. This book won't just teach you how to work with Exchange databases, transport services, mail flow, and Client Access servers; it will teach you how to squeeze every last bit of power out of these features and options while making the most of what Exchange Server 2013 provides.

Unlike many other books about administering Exchange Server 2013, this book doesn't focus on a specific user level. This isn't a lightweight beginner book. Regardless of whether you are a beginning administrator or a seasoned professional, many of the concepts in this book will be valuable to you, and you can apply them to your Exchange Server 2013 installations.

Who is this book for?

Microsoft Exchange Server 2013 Pocket Consultant: Databases, Services, & Management covers the Standard and Enterprise editions of Exchange Server 2013. The book is designed for the following readers:

- Current Exchange Server 2013 administrators
- Current Windows administrators who want to learn Exchange Server 2013
- Administrators upgrading to Exchange Server 2013 from Exchange Server 2007 or Exchange Server 2010
- Administrators transitioning to Exchange Server 2013 from Exchange Server 2003
- Administrators transferring from other messaging servers
- Managers and supervisors who have been delegated authority to manage mailboxes or other aspects of Exchange Server 2013

To pack in as much information as possible, I had to assume that you have basic networking skills and a basic understanding of email and messaging servers. With this in mind, I don't devote entire chapters to explaining why email systems are

needed or how they work, nor do I devote entire chapters to installing Exchange Server 2013. I do, however, provide complete details on the components of Exchange organizations and how you can use these components to build a fully redundant and highly available messaging environment. You will also find complete details on all the essential Exchange administration tasks for availability groups, Exchange databases, mail flow, transport services, Client Access servers, and much more.

I also assume that you are fairly familiar with Windows Server. If you need help learning Windows Server, I highly recommend that you buy *Windows Server 2012 Pocket Consultant* or *Windows Server 2012 Inside Out*.

How is this book organized?

Rome wasn't built in a day, nor was this book intended to be read in a day, or in a week, or even in a month for that matter. Ideally, you'll read this book at your own pace, a little each day as you work your way through each of the nine chapters. The chapters are arranged in a logical order, taking you from planning for availability groups and databases to Exchange Server maintenance and disaster recovery.

Ease of reference is an essential part of this hands-on guide. This book has an expanded table of contents and an extensive index for finding answers to problems quickly. Many other quick-reference features have been added to the book as well, including quick step-by-step procedures, lists, tables with fast facts, and extensive cross references.

As with all Pocket Consultants, *Microsoft Exchange Server 2013 Pocket Consultant: Databases, Services, & Management* is designed to be a concise and easy-to-use resource for managing Exchange servers. This is the readable resource guide that you'll want on your desktop at all times. The book covers everything you need to perform the core administration tasks for the following:

- Managing availability groups and Exchange databases
- Managing mail flow and transport services
- Working with Client Access servers
- Managing mobile messaging users
- Maintaining and monitoring Exchange servers
- Backing up and restoring Exchange servers

Although designed and written to stand on its own, this book also can be used with *Microsoft Exchange Server 2013 Pocket Consultant: Configuration & Clients*, which focuses on the following:

- Deploying Exchange Server 2013
- Exchange administration essentials
- Managing Exchange clients
- Administration of users, contacts, and mailboxes
- Configuring distribution groups and address lists
- Implementing Exchange Server security and permissions

Because the focus is on giving you maximum value in a pocket-size guide, you don't have to wade through hundreds of pages of extraneous information to find what you're looking for. Instead, you'll find exactly what you need to get the job done, and you'll find it quickly.

Conventions used in this book

I've used a variety of elements to help keep the text clear and easy to follow. You'll find code terms and listings in monospace type, except when I tell you to actually enter a command; in which case, the command appears in bold type. When I introduce and define a new term, I put it in italics.

Other conventions include the following:

- **Best practices** To examine the best technique to use when working with advanced configuration and administration concepts
- **Caution** To warn you of potential problems
- **Important** To highlight important concepts and issues
- **More info** To provide more information on the subject
- **Note** To provide details on a point that needs emphasis
- **Real world** To provide real-world advice when discussing advanced topics
- **Security alert** To point out important security issues
- **Tip** To offer helpful hints or additional information

I truly hope you find that *Microsoft Exchange Server 2013 Pocket Consultant: Databases, Services, & Management* provides everything you need to perform essential administrative tasks as quickly and efficiently as possible. You are welcome to send your thoughts to me at *williamstanek@aol.com*. Follow me on Twitter at WilliamStanek and on Facebook at *www.facebook.com/William.Stanek.Author*.

Other resources

No single magic bullet for learning everything you'll ever need to know about Exchange Server 2013 exists. Although some books are offered as all-in-one guides, there's simply no way one book can do it all. With this in mind, I hope you use this book as it is intended to be used—as a concise and easy-to-use resource. It covers everything you need to perform core administration tasks for availability groups, databases, transport services, mail flow, and Client Access servers, but it is by no means exhaustive.

Your current knowledge will largely determine your success with this or any other Exchange resource or book. As you encounter new topics, take the time to practice what you've learned and read about. Seek out further information as necessary to get the practical hands-on know-how and knowledge you need.

For topics this book doesn't cover, you might want to look to *Microsoft Exchange Server 2013 Pocket Consultant: Configuration & Clients*. I also recommend that you regularly visit the Microsoft website for Exchange Server (*microsoft.com/exchangeserver/*) and *support.microsoft.com* to stay current with the latest changes. To help you get the most out of this book, you can visit my corresponding website at *pocket-consultant.com*. This site contains information about Exchange Server 2013 and updates to the book.

Errata and book support

Every effort has been made to ensure the accuracy of this book and its companion content. Any errors that have been reported since this book was published are listed on our Microsoft Press site:

http://aka.ms/ExDSM/errata

If you find an error that is not already listed, you can report it to us through the same page.

If you need additional support, email Microsoft Press Book Support at *mspinput@microsoft.com*.

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

We want to hear from you

At Microsoft Press, your satisfaction is our top priority, and your feedback is our most valuable asset. Please tell us what you think of this book at:

http://www.microsoft.com/learning/booksurvey

The survey is short, and we read every one of your comments and ideas. Thanks in advance for your input!

Stay in touch

Let's keep the conversation going! We're on Twitter: *http://twitter.com/MicrosoftPress*.

CHAPTER 3

Exchange database administration

- Working with active mailbox databases **63**
- Working with mailbox database copies **78**
- Managing mailbox databases **95**
- Content indexing **103**

Databases are containers for information. Microsoft Exchange Server 2013 uses mailbox databases to maintain user data. The information in a particular database isn't exclusive to mailboxes and their associated user data—Exchange Server maintains related information within databases as well. Within mailbox databases, you'll find information about Exchange logons and mailbox usage. Exchange also maintains information about full-text indexing, although the actual content indexes are stored in separate files. In this chapter, you'll learn how to manage databases and the information they contain.

Working with active mailbox databases

Each Mailbox server installed in the organization has an information store. The information store operates as a service and manages the server's databases. Each mailbox database has a database file associated with it. This file is stored in a location that you specify when you create or modify the mailbox database.

Mailbox databases can be either active databases or passive copies of databases. Users access active databases to get their mailbox data. Passive copies of databases are not actively being used and are the subject of the section, "Working with mailbox database copies" later in this chapter. You create passive copies of databases as part of a high-availability configuration as discussed in Chapter 2, "Managing data and availability groups."

Understanding mailbox databases

Mailboxes are the normal delivery location for messages delivered within an organization. They contain messages, attachments, and other types of information that the user might have placed in the mailbox. Mailboxes, in turn, are stored in mailbox databases.

When you install a Mailbox server, Setup creates a default mailbox database. The default mailbox database is meant to be a starting point, and most Exchange organizations can benefit from having additional mailbox databases, especially as the number of users in the organization grows. Additional mailbox databases are created for many reasons, but the following reasons are the most common:

- **To provide a smaller unit of management** Exchange has a practical limit of 2 terabytes (TB) on the size of databases, though you may find it easier to work with databases between 1 TB and 1.5 TB. Large databases require more time to move, restore, and recover compared to smaller databases. Additionally, when you establish database availability groups and create copies of a database, the entire database must be replicated from the source database to the database copies. During recovery, you can restore individual databases without affecting the performance or uptime of other databases on the system.
- **To impose a different set of mailbox rules on different sets of users** Each additional mailbox database can have its own property settings for maintenance, storage limits, deleted item retention, indexing, security, and policies. By placing a user's mailbox in one mailbox database instead of another, you can apply a different set of rules.
- **To optimize Exchange performance** Each mailbox database can have its own storage location. By placing the mailbox databases on different physical drives, you can improve the performance of Exchange Server 2013.
- **To create separate mailbox databases for different purposes** For example, you might want to create a mailbox database called General In-Out to handle all general-purpose mailboxes being used throughout the organization. These general-purpose mailboxes could be set up as shared mailboxes for Postmaster, Webmaster, Technical Support, Customer Support, and other key functions.

When you create a mailbox database, you can specify the following information:

- What the name of the database should be
- Where the database file is to be located
- When maintenance on the database should occur
- Any limitations on mailbox size
- Whether deleted items and mailboxes should be retained

Each mailbox database has a default offline address book (OAB). Microsoft Outlook 2007 and later clients access the default OAB and default public folder hierarchy on your organization's Client Access servers. Exchange 2013 uses the mailbox provisioning load balancer to automatically select a database to use when you create or move a mailbox and do not explicitly specify the mailbox database to use. As the

name implies, the purpose of the load balancer is to try to balance the workload across mailbox databases in the organization.

Although the load balancer uses multiple criteria to try to determine where a mailbox should be created or moved, the selection criteria does not take into account the proximity of the Mailbox server on which a database is stored to the computer or computers used by the user. Instead, the load balancer uses the Active Directory site where the mailbox task is being performed to determine which mailbox databases should be selected and only includes databases that are in the local site.

You can control the way automatic distribution works in several ways. You can temporarily or permanently exclude databases from the distribution process by using the -IsSuspendedFromProvisioning and -IsExcludedFromProvisioning parameters of the Set-MailboxDatabase cmdlet respectively. When either of these parameters is set to $True, Exchange excludes the related database from the automatic distribution process.

When selecting a database to use, the mailbox provisioning load balancer also checks the database management scopes of the administrator creating a mailbox. Database management scopes are part of the role-based access control (RBAC) permissions model and are a way to limit the databases administrators can view and manage.

> ***NOTE*** **By default, all administrators in an Exchange organization can see all the mailbox databases in the organization. When you create database management scopes in the organization, administrators will only be able to see databases included in a scope applied to them.**

If you create custom scopes, Exchange uses these scopes to select databases. Specifically, the load balancer only selects mailbox databases included in a scope applied to the administrator creating a mailbox. Therefore, if a database isn't included in a scope applied to an administrator, the database won't be selected for automatic distribution.

Preparing for automatic reseed

Automatic reseed, new for Exchange 2013, allows you to quickly restore database redundancy after a disk failure, database corruption event, or other event that requires a reseed of a database to recover operations. For automatic reseed to work, however, you must pre-provision one or more spare disks. These spare disks are then used during the automatic reseed to recover the database copy. Here's how automatic reseed works:

1. The Microsoft Exchange Replication service scans the Information Store periodically for database copies that have a status of FailedAndSuspended.
2. If the replication service finds a database copy with the FailedAndSuspended status, it performs prerequisite checks to evaluate the situation, which includes determining whether spares are available, whether anything could prevent the system from performing an automatic reseed, whether only a single copy of the database is available, and more.

3. If the prerequisite checks pass successfully, the Microsoft Exchange Replication service allocates and remaps an available spare before starting the seed operation.
4. After the seed has been completed, the Microsoft Exchange Replication service verifies that the new copy has a Healthy status.

To prepare spare volumes on a server, you must complete the following steps:

1. Mount the volumes that will contain databases under a single mount point, such as C:\PrimaryVols.
2. Mount the volumes to mount points under this volume. For example, you could mount the first volume as C:\PrimaryVols\Volume1, the second volume as C:\PrimaryVols\Volume2, and so on.
3. Create databases on the server in locations within the specified volumes, ensuring that there are fewer databases than mounted volumes.

Consider the following scenario to see how this would work in practice:

You have five volumes mounted under C:\PrimaryVols as C:\PrimaryVols\Volume1, C:\PrimaryVols\Volume2, C:\PrimaryVols\Volume3, C:\PrimaryVols\Volume4, and C:\PrimaryVols\Volume5.

You create three databases, locating the first database under C:\PrimaryVols\Volume1, the second under C:\PrimaryVols\Volume2, and the third under C:\PrimaryVols\Volume3.

You then have two spare volumes, mounted as C:\PrimaryVols\Volume4 and C:\PrimaryVols\Volume5.

If a disk fails, a database copy becomes corrupted, or another event requiring reseed occurs, the failed database is automatically reseeded to one of the spare volumes.

You can identify the failure and automatic reseed tasks by reviewing the event logs. Related events are logged in the event logs under Applications and Services Logs > Microsoft > Exchange > High Availability and under Applications and Services Logs > Microsoft > Exchange >MailboxDatabaseFailureItems.

Creating mailbox databases

You can create mailbox databases by using the New Mailbox Database Wizard. The default database file path and default log folder path are set automatically to be the same as those used for other Exchange data.

Any new mailbox databases you create using the Exchange Admin Center are configured to use the mailbox provisioning load balancer by default. When you create mailbox databases using the Exchange Management Shell, you can set the –IsExcludedFromProvisioning parameter to $True to specify that the database should not be considered by the mailbox provisioning load balancer. Excluding a database from provisioning means new mailboxes are not automatically added to this database. Rather than excluding a database from provisioning, you can set the –IsSuspendedFromProvisioning parameter to $True to specify that a database temporarily not be considered by the mailbox provisioning load balancer. Keep in mind

that whether you exclude or suspend a database from provisioning is semantics as in either case the database won't be used for provisioning.

To create a mailbox database, complete the following steps:

1. In the Exchange Admin Center, select Servers in the feature pane, and then select Databases. In the main pane, you should see a list of active databases that are available in the Exchange organization.
2. Tap or click the New button to open the New Database dialog box, shown in Figure 3-1.

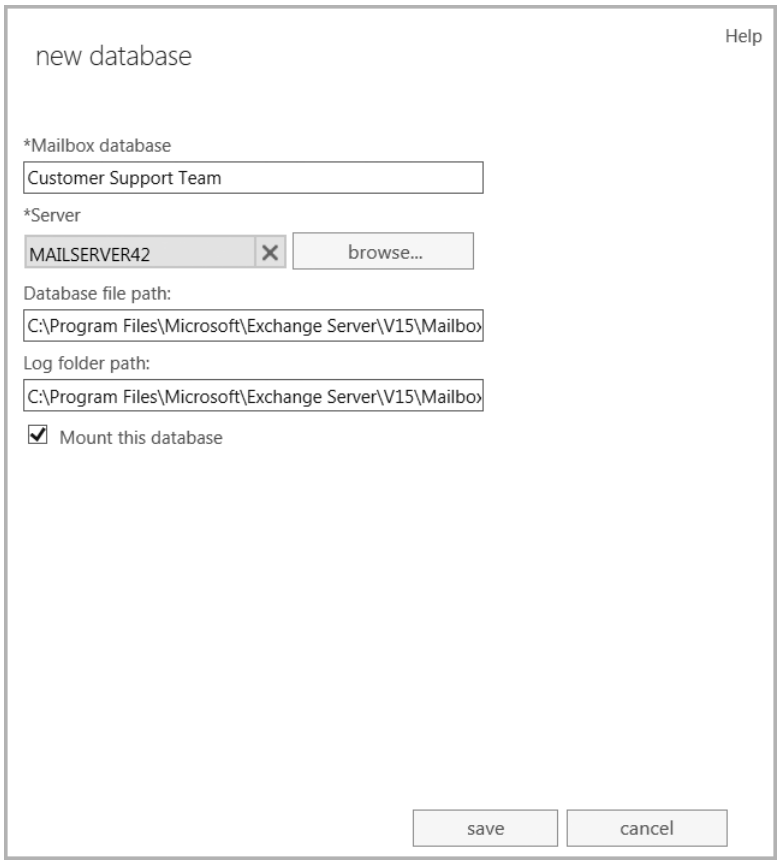

FIGURE 3-1 Configure the new mailbox database.

3. In the Mailbox Database text box, type a name for the mailbox database. The database name must be unique within the Exchange organization. Although the database name can contain spaces and special characters, it'll be easier to work with the database if the name uses standard characters.
4. Tap or click Browse to the right of the Server text box to open the Select Server dialog box. Figure 3-2 shows Mailbox servers listed by name, version, and exact build as well as associated database availability group, if applicable.

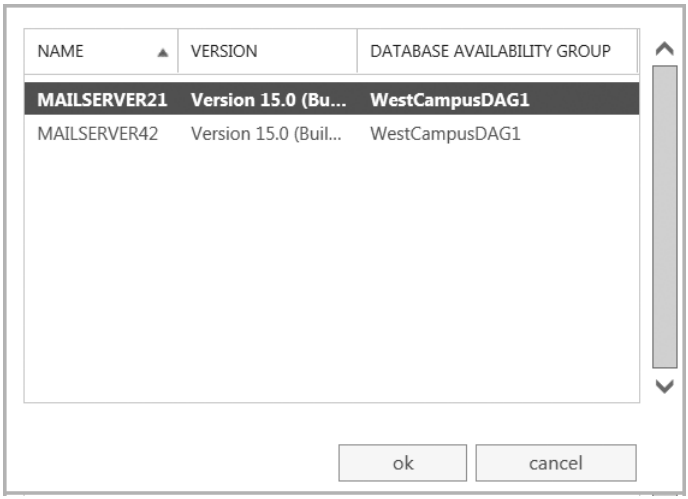

FIGURE 3-2 Select a Mailbox server.

5. Select the Mailbox server that will host the mailbox database, and then tap or click OK. Only Mailbox servers in the Active Directory forest to which you are connected are available.
6. The database file path and log folder path are set to the default location for Exchange data on the selected server. A subfolder with the mailbox database name will be created under the default database file path and the name of the .edb file for the database will be set the same as the database name. Similarly, a subfolder with the same name as the database name is created under the default log folder path. If you don't want to use the default locations, enter the paths you want to use for the database file and the related logs in the text boxes provided.

REAL WORLD **Exchange creates folders if they do not exist, which is a good thing except when you mistype the intended path. Rather than type in a long file path, you might want to use copy and paste. In File Explorer, navigate to the exact folder path you want to use. Click in the folder path on the Address bar to display and automatically select the folder path. Press Ctrl+C to copy the path. In the New Database dialog box, click in the path text box, press Ctrl+A, and then press the Delete key. Finally, press Ctrl+V to paste in the path you copied previously.**

7. Select the Mount This Database check box if you want to mount this database. Mounting a database puts it online, making it available for use.
8. Tap or click Save to create the mailbox database, and then tap or click OK. If an error occurred, you need to take the appropriate corrective action. Otherwise, you can now modify the properties of the mailbox database as necessary. To make the new database accessible to mailbox users, you must restart the Microsoft Exchange Information Store service.

Exchange Server 2013 Standard edition supports up to five databases. Exchange Server 2013 Enterprise edition supports up to 100 databases. However, if you install Exchange Server 2013 Enterprise edition but forget to enter the product key, the server runs in Trial mode and only supports up to five databases as well.

If you try to create more databases than are supported by your edition of Exchange, you'll see an error stating that Exchange couldn't mount the database. In the error details, you'll also see a message stating MapiExceptionTooManyMounted-Databases: Unable To Mount Database (see Figure 3-3). You can resolve this problem on a server running Enterprise edition by reducing the number of databases on the server or simply creating the database on a different server. To resolve this problem on a server running a standard or trial edition of Exchange Server 2013, you can upgrade the server to Enterprise edition by completing these steps:

1. In Exchange Admin Center, select Servers on the feature pane, and then select Servers.
2. Double-tap or double-click the server you want to upgrade. In the Properties dialog box, on the General page, the current edition should be listed as Trial Edition or Standard Edition.

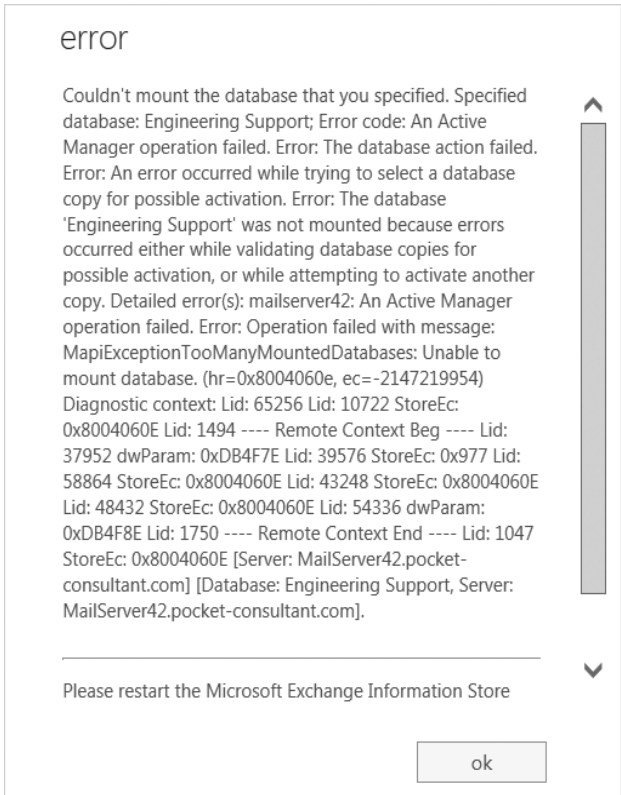

FIGURE 3-3 An error showing that Exchange Admin Center was unable to mount the new database.

3. If the server is running Trial Edition, upgrade by entering a valid Enterprise product key in the text boxes provided, and then selecting Save. If the server is running Standard Edition, upgrade by selecting Change Product Key, entering a valid Enterprise product key in the text boxes provided, and then selecting Save.
4. Tap or click OK. For the change to take effect, you must restart the Microsoft Exchange Information Store service.

In the Exchange Management Shell, you can create mailbox databases by using the New-MailboxDatabase cmdlet. Listing 3-1 provides the syntax and usage.

NOTE **You use a separate cmdlet to mount the database. See the section "Mounting and dismounting databases" later in this chapter for details.**

LISTING 3-1 New-MailboxDatabase cmdlet syntax and usage

Syntax

```
New-MailboxDatabase -Name DatabaseName -Server ServerName
[-EdbFilePath DbFilePath] [-LogFolderPath FolderPath] {AddtlParams}

{AddtlParams}
[-DomainController FullyQualifiedName][-IsExcludedFromProvisioning <$true
| $false>] [-IsSuspendedFromProvisioning <$true | $false>]
[-OfflineAddressBook OfflineAddressBook]


New-MailboxDatabase -Recovery <$true | $false> -Server ServerName
[-DomainController FullyQualifiedName] [-EdbFilePath DbFilePath]
[-LogFolderPath FolderPath]
```

Usage

```
New-MailboxDatabase –Server "CorpServer88" –Name "Accounting Database"
-EdbFilePath "C:\Databases\Accounting\AccountingMail.edb"
-LogFolderPath "D:\DatabaseLogs\Accounting"
```

Setting the default offline address book

Mailbox databases can have different types of information associated with them, including a default OAB. You set related options for mailbox databases by using the Client Settings page of the related Properties dialog box. To view this dialog box and update the messaging options, follow these steps:

1. In the Exchange Admin Center, select the Servers feature, and then select Databases. Next, double-tap or double-click the database you want to configure.
2. In the Properties dialog box, tap or click the Client Settings page.

NOTE **If you can't update the text boxes on the Client Settings page, it means that a policy has been applied to the mailbox database. You must directly edit or remove the policy and then make the necessary changes.**

3. The Offline Address Book text box shows the OAB for the mailbox database. OABs contain information regarding mail-enabled users, contacts, and groups in the organization, and they are used when users aren't connected to the network. If the text box is empty, the global default is used. If you've created additional OABs beyond the global default, you can specify one of these additional OABs as the default for the mailbox database. Tap or click Browse, select the OAB you want to use, and then tap or click OK. Tap or click Save to apply the changes.

In the Exchange Management Shell, you can set the default OAB for mailbox databases by using the Set-MailboxDatabase cmdlet. Listing 3-2 provides the syntax and usage.

LISTING 3-2 Using the Set-MailboxDatabase cmdlet to set the default OAB

Syntax

```
Set-MailboxDatabase –Identity MailboxDatabase
[-OfflineAddressBook OABIdentity
```

Usage

```
Set-MailboxDatabase –Identity "Accounting Mail"
-OfflineAddressBook "\US Corporate"
```

Setting mailbox database limits and deletion retention

Mailbox database limits are designed to control the amount of information that users can store in their mailboxes. Users who exceed the designated limits might receive warning messages and might be subject to certain restrictions, such as the inability to send messages. Deleted item retention is designed to ensure that messages and mailboxes that might be needed in the future aren't deleted inadvertently. If retention is turned on, you can retain deleted messages and mailboxes for a specified period before they are permanently deleted and are nonrecoverable.

An average retention period for messages is about 14 days. The minimum retention period for mailboxes should be about 7 days. In most cases, you'll want deleted messages to be maintained for a minimum of 5 to 7 days and deleted mailboxes to be maintained for a minimum of three to four weeks. An interval of 5 to 7 days is used for messages because users usually realize within a few days that they shouldn't have deleted a message. A three-week to four-week interval is used for mailboxes because several weeks can (and often do) pass before users realize that they need a deleted mailbox or messages within a deleted mailbox. To understand why, consider the following scenario.

Sally leaves the company. A coworker is given permission to delete Sally's user account and mailbox. Three weeks later, Sally's boss realizes that she was the only person who received and archived the monthly reports sent through email from corporate headquarters. The only way to get reports for previous years is to recover Sally's mailbox, and you can do this if you've set a sufficiently long retention period.

NOTE **Exchange has several features to ensure that mailbox items are retained according to policies set forth by an organization for legal reasons, including automatic archiving of old messages and retention policies. Deletion settings on the Limits page control the minimum length of time deleted items are retained if no retention tags specifically apply to deleted items.**

To view or set limits and deletion retention for a mailbox database, follow these steps:

1. In the Exchange Admin Center, select the Servers feature, and then select Databases. Next, double-tap or double-click the database you want to configure.
2. In the Properties dialog box, on the Limits page (shown in Figure 3-4), use the following options to set storage limits and deleted item retention:
   - **Issue a warning at (GB)** Sets the size limit, in gigabytes, that a mailbox can reach before Exchange Server issues a warning to the user. The warning tells the user to clear out the mailbox. If you don't want Exchange to issue warnings, set the value to 0 or Unlimited.
   - **Prohibit send at (GB)** Sets the size limit, in gigabytes, that a mailbox can reach before the user is prohibited from sending any new mail. The restriction ends when the user clears out the mailbox and the total mailbox size is under the limit. If you don't want Exchange to prohibit sending mail, set the value to 0 or Unlimited.
   - **Prohibit send and receive at (GB)** Sets the size limit, in gigabytes, that a mailbox can reach before the user is prohibited from sending and receiving mail. The restriction ends when the user clears out the mailbox and the total mailbox size is under the limit. If you don't want Exchange to prohibit sending and receiving mail, set the value to 0 or Unlimited.

CAUTION **Prohibiting send and receive might cause users to lose email. When a user sends a message to a user who is prohibited from receiving messages, a non-delivery report (NDR) is generated and delivered to the sender. The recipient never sees the email. Because of this, you should prohibit send and receive only in very rare circumstances. Your organizational policy will likely spell out those circumstances. To remove this restriction, set Prohibit Send And Receive to Unlimited or enter a value of 0.**

   - **Keep deleted items for (days)** Sets the number of days to retain deleted items. An average retention period is 14 days. If you set the retention period to 0, deleted messages aren't retained, and you can't recover them in the same way you could if retention was enabled.

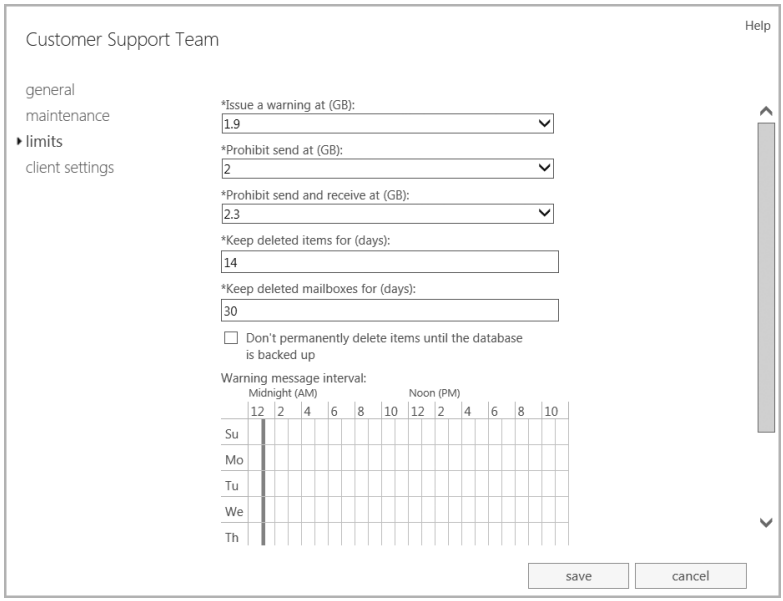

FIGURE 3-4 Use the Limits page to set storage limits and deleted item retention for individual mailboxes and entire mailbox databases.

- **Keep deleted mailboxes for (days)** Sets the number of days to retain deleted mailboxes. The default setting is 30 days. You'll want to keep most deleted mailboxes for at least 7 days to allow the administrators to extract any data that might be needed. If you set the retention period to 0, deleted mailboxes are retained only if you select the next option, and then only until the database has been backed up. If a mailbox is backed up, you can recover it only by restoring it from backups.
- **Don't permanently delete items until the database is backed up** Ensures that deleted mailboxes and items are archived into at least one backup set before they are removed.

3. The Warning Message Interval sets the interval for sending warning messages to users whose mailboxes exceed the designated limits. To change this setting, select Customize. You can now set the warning interval using the Customize Quota Notification Schedule dialog box, shown in Figure 3-5.
   - Times that are used for quota notification are filled in with a dark bar.
   - Times that aren't used for quota notification are blank.

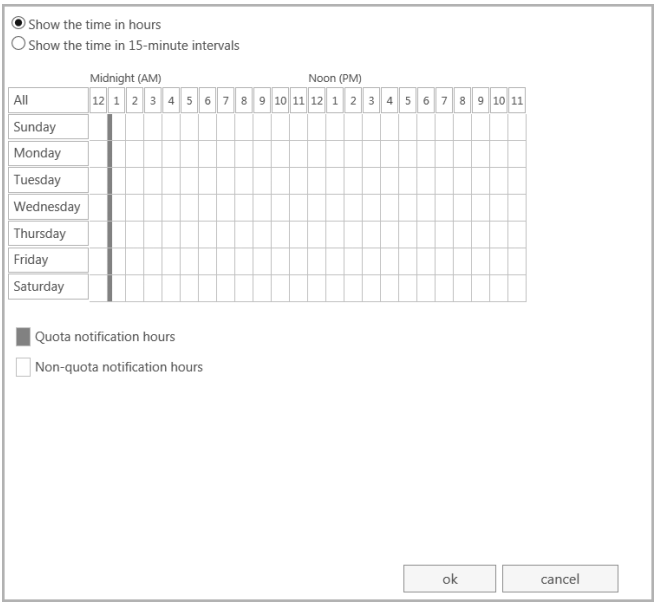

FIGURE 3-5 Customize quota notification for mailboxes stored in the database.

IMPORTANT **The default interval for sending warning messages is daily between 1 A.M. and 1:15 A.M, which is an acceptable initial interval for small deployments. As your organization grows, however, you'll want to optimize this interval to ensure that servers aren't overburdened and that servers have enough time to process all the mailboxes.**

4. Show the time in hours or in 15-minute intervals by using the options provided. Click or tap the time interval to change the setting.
   - Hourly or 15-minute interval buttons are used to select or clear a particular interval for all the days of a week.
   - Days of the week buttons allow you to clear or select all the hours in a particular day.
   - The All button allows you to clear or select all the time periods.
5. If you customized the notification schedule, tap or click OK to close the Customize Quota Notification Schedule dialog box.
6. Tap or click Save to apply your settings.

In the Exchange Management Shell, you can set limits for mailbox databases by using the Set-MailboxDatabase cmdlet. Listing 3-3 provides the syntax and usage. When you set a limit, you can specify the value with KB (for kilobytes), MB (for megabytes), or GB (for gigabytes). The default value type is bytes. Additionally, it's important to point out that the -MaintenanceSchedule and -QuotaNotificationSchedule parameters are not used with Exchange 2013.

LISTING 3-3 Using the Set-MailboxDatabase cmdlet to set limits

Syntax

```
Set-MailboxDatabase [-Identity MailboxDatabase]
[-AllowFileRestore <$true | $false>] [-BackgroundDatabaseMaintenence <$true
| $false>] [-CircularLoggingEnabled <$true | $false>]
[-DataMoveReplicationConstraint <None | SecondyCopy | SecondDatacenter |
AllDatacenters | AllCopies>] [-DeletedItemRetention NumberDays]
[-DomainController DCName] [-EventHistoryRetentionPeriod NumberDays]
[-IndexEnabled <$true | $false>] [-IsExcludedFromProvisioning <$true |
$false>] [-IssueWarningQuota Limit] [-JournalRecipient RecipientId]
[-MailboxRetention NumberDays] [-MountAtStartup <$true | $false>]
[-Name Name] [-OfflineAddressBook OABId] [-ProhibitSendQuota Limit]
[-ProhibitSendReceiveQuota Limit] [-RecoverableItemsQuota Limit]
[-RecoverableItemsWarningQuota Limit]
[-RetainDeletedItemsUntilBackup <$true | $false>]
```

Usage

```
Set-MailboxDatabase –Identity "Accounting Mail"
 -IssueWarningQuota 1.9GB
 -DeletedItemRetention 14
 -MailboxRetention 30
 -ProhibitSendQuota 2GB
 -ProhibitSendReceiveQuota 2.4GB
 -RetainDeletedItemsUntilBackup $true
```

Recovering deleted mailboxes

When you delete a mailbox from a user account, the mailbox is retained as a disconnected mailbox according to the mailbox retention setting. You can reconnect the mailbox to the original user account or another user account if necessary. Similarly, when you delete a user account and the related mailbox, the mailbox is retained as a disconnected mailbox according to the mailbox retention setting. You can connect the mailbox to an existing user account if necessary.

When you move mailboxes between databases, mailboxes in the original (source) database are soft deleted. This means they are disconnected, marked as soft deleted, but retained in the original database until the deleted mailbox retention period expires. In Exchange Management Shell, you can use a DisconnectReason of "Soft-Deleted" to find soft-deleted mailboxes.

To recover a deleted mailbox, complete the following steps:

1. In the Exchange Admin Center, select Recipients in the feature pane, and then select Mailboxes.
2. Tap or click the More button (this button shows three dots), and then select Connect A Mailbox. The Connect A Mailbox dialog box shows all mailboxes marked for deletion but currently retained regardless of whether those mailboxes were disabled, deleted, or soft deleted.

3. In the Connect A Mailbox dialog box, shown in Figure 3-6, use the selection list provided to select the server in which you want to look for disconnected mailboxes.

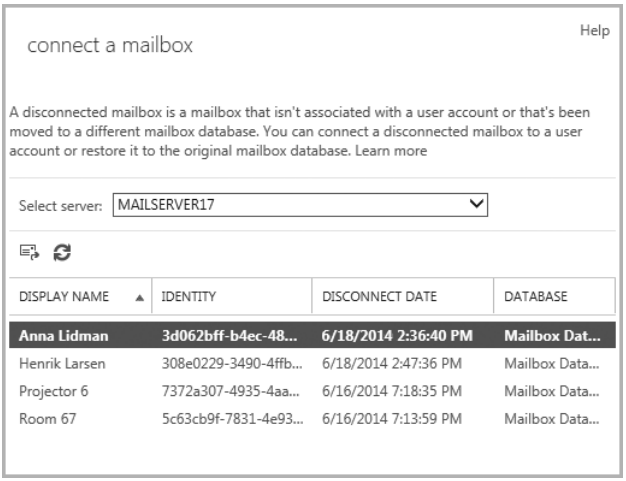

FIGURE 3-6 Viewing disconnected mailboxes.

4. Tap or click the mailbox to restore it, and then tap or click Connect. Connect the mailbox to the user account to which it was connected previously or to a different user account. If the original user account is available, select the Yes option to reconnect the mailbox to the original user account. If the original user isn't available or you want to associate the mailbox with a different user, select the No option and follow the prompts.

NOTE **Deleted mailboxes aren't necessarily marked as such immediately. It can take 15 minutes to an hour before the mailbox is marked as deleted and listed accordingly.**

IMPORTANT **If you previously removed the mailbox rather than disabling it, the user account associated with the mailbox was deleted as well. Because each user account has a unique security identifier associated with it, you can't simply re-create the user account to get back the same set of permissions and privileges. That said, if you want to connect the mailbox to a user account with the same name, you can do this by recovering the deleted account from Active Directory before garbage collection has occurred or by recreating the account in Active Directory Users And Computers. The account will then be available for selection but when you're connecting the mailbox to an account, you'll need to choose the No option because Exchange and Active Directory see this as a different account.**

You can use the Connect-Mailbox cmdlet to perform the same task following the syntax shown in Listing 3-4.

LISTING 3-4 Connect-Mailbox cmdlet syntax and usage

Syntax

```
Connect-Mailbox -Identity OrigMailboxIdentity
-Database DatabaseIdentity
-User NewUserIdentity
[-ActiveSyncMailboxPolicy PolicyId] [-Alias Alias]
[-DomainController DCName] [-ManagedFolderMailboxPolicy PolicyId]
[-ManagedFolderMailboxPolicyAllowed <$true | $false>]
[-Archive <$true | $false>] [-Equipment <$true | $false>]
[-Room <$true | $false>] [-Shared <$true | $false>]
[-ValidateOnly <$true | $false>]

[-LinkedCredential Credential] [-LinkedDomainController DCName]
[-LinkedMasterAccount UserId]
```

Usage

```
Connect-Mailbox -Identity "Per Reitzel"
-Database "Accounting Mail" -User "CPANDL\perr" -Alias "perr"

Connect-Mailbox -Identity "Per Reitzel"
-Database "Accounting Mail" –LinkedDomainController CorpServer72
-LinkedMasterAccount "CPANDL\perr"
```

Recovering deleted items from mailbox databases

You can recover deleted items from mailbox databases as long as you've either set a deleted item retention period for the database from which the items were deleted and the retention period hasn't expired, or you have specified that Exchange should not permanently delete items from mailboxes until the database has been backed up and Exchange hasn't been backed up yet. If either of these conditions are met, you can recover deleted items from mailbox databases.

To use Outlook 2010 or Outlook 2013 for recovery, complete the following steps:

1. Log on as the user who deleted the message, and then start Outlook.
2. Tap or click the Folders pane, and then select Recover Deleted Items.
3. The Recover Deleted Items From dialog box appears. Select the items you want to recover, and then tap or click the Recover Selected Items button.
4. Items you've recovered are copied to the Deleted Items folder. In the left pane, tap or click Deleted Items.
5. In the Deleted Items folder, press and hold or right-click items you want to keep, select Move, and then tap or click Other Folder.
6. In the Move Items dialog box, select the folder to which the item should be moved, and then tap or click OK.

NOTE **The steps are similar for Outlook 2007, except you start by tapping or clicking Recover Deleted Items on the Tools menu.**

To use Outlook Web App (OWA) for recovery, complete these steps:

1. In a web browser, type ***https://servername.yourdomain.com/owa***, where *servername* is a placeholder for the HTTP virtual server hosted by Exchange Server 2013 and *yourdomain.com* is a placeholder for your external domain name, such as https://mail.cpandl.com/owa.
2. Next, log on as the user (or have the user log on). Type the user name in *domain\username* format, such as **pocket-consulta\bertk,** or *user@domain* format, such as **berk@pocket-consultant.com**. Type the password, and then tap or click Sign In.
3. In the left pane, press and hold or right-click Deleted Items, and then select Recover Deleted Items.
4. In the Recover Deleted Items dialog box, you'll see a list of recoverable items. Each listed item will have a selection check box. Select this check box for items you want to recover.
5. Tap or click the Recover button, and then tap or click OK. Items you select will be restored to their default folders.

Working with mailbox database copies

When your Exchange organization uses database availability groups, Exchange replicates transaction logs from an active mailbox database on a source Mailbox server to other Mailbox servers in the database availability group that have passive copies of the database. On these servers, Exchange replays the transaction logs into the passive copy of the mailbox database by using either file mode or block mode replication. You can monitor the health and status of replication and database copies by using the Exchange Management tools.

The Mailbox server that hosts the active copy of a database is referred to as the *mailbox database primary* for that database. A Mailbox server that hosts a passive copy of a database is referred to as a *mailbox database secondary* for that database. You can move the active database to another Mailbox server in the database availability group by using the switchover process discussed in "Switching over servers and databases" in Chapter 2. In a switchover, the active copy of a database is dismounted on the current Mailbox server and a passive copy of the database is activated and mounted on another Mailbox server in the database availability group.

TIP **You can quickly distinguish between an active mailbox database and a passive copy of a database by reviewing the Copy Status column under the Database Copies page in the Exchange Admin Center. Only the active database will have a status of Mounted or Dismounted. For passive database copies, you'll see the current status of replication for the database copy.**

Creating mailbox database copies

After you create a database availability group and add Mailbox servers to the group, you can create copies of mailbox databases to initiate replication. Within the group, replication occurs between the active mailbox database on a source Mailbox server and other Mailbox servers that host copies of the database. You cannot replicate a database outside of a database availability group, nor can you replicate an Exchange 2013 mailbox database to a server running an earlier version of Exchange.

Each database availability group can have up to 16 member servers, and you can create up to 16 instances of a database, including one active instance and 15 passive instances. You can create mailbox database copies only on Mailbox servers that do not host the active copy of a mailbox database, and you cannot create two copies of the same database on the same server.

Because all copies of a database use the same path on each server containing a copy, the database and log file paths for a database copy on each Mailbox server must not conflict with any other database paths. You need to ensure the database and log file paths for the database copy can be created in the same location as all other copies and that the paths do not conflict with any other database paths on the target server.

With respect to Active Directory, the member servers in an availability group must all be in the same Active Directory domain. You can create database copies on Mailbox servers in the same or different Active Directory Sites, and on the same or different network subnets. However, database copies are not supported between Mailbox servers with roundtrip network latency greater than 250 milliseconds (by default). Database copies are automatically assigned an identity in the format *DatabaseName**HostMailboxServerName*, such as Engineering Primary Database\MailServer36.

To create a copy of a mailbox database, complete the following steps:

1. In the Exchange Admin Center, select Servers in the feature pane, and then select Databases. In the main pane, you should see a list of active databases that are available in the Exchange organization.
2. Select the mailbox database that you want to copy to see a list of all copies of that database in the details pane. Whereas the active copy of a database normally is listed with a status of Active Mounted or Active Dismounted, passive copies are normally listed with a status of Passive Healthy.
3. Tap or click the More button (this button shows three dots), and then select Add A Database Copy. This opens the Add Mailbox Database Copy dialog box.

   As shown in Figure 3-7, the dialog box shows you which servers already have a copy of the database and sets the activation preference number to the next value for the next database instance. You can set a lower preference value if desired.

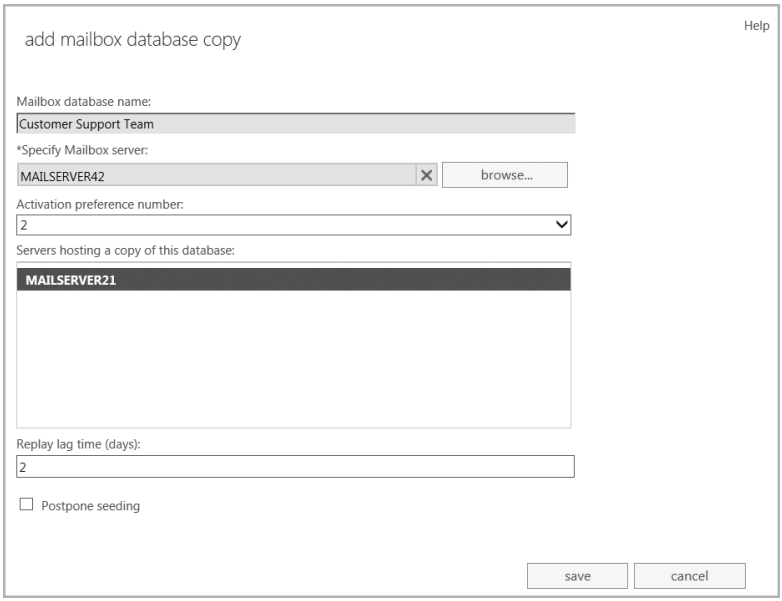

FIGURE 3-7 Add a mailbox database copy.

4. Tap or click Browse. Select the Mailbox server that will host the mailbox database copy, and then tap or click OK. Although servers outside the database availability group and servers running earlier versions of Exchange may be listed in the Select Server dialog box, you'll only want to select an Exchange 2013 Mailbox server in the same database availability group that doesn't have a copy of the database already. Each Mailbox server in a database availability group can host only one copy of a database.
5. Optionally, in the Activation Preference Number text box, specify the preference value for the database copy. The activation preference number represents the order of activation preference for a database copy after a failure or outage of the active copy. The preference value is a number equal to or greater than 1, where 1 has the highest preference. The preference value cannot be larger than the total number of database copies.

NOTE **Active Manager uses the preference value only to break ties in the best-copy selection process. If two or more database copies are seen as the best choice for activation, the database copy with the highest preference is selected. Following this, when there is a tie, a database copy with a preference value of 3 would be selected before a database copy with a preference value of 4. For more information on Active Manager, see "Working with active manager" in Chapter 2.**

6. If you want to configure replay lag time or postpone seeding, select More Options. You'll then be able to specify a replay lag time in days and postpone seeding. If you postpone seeding of the database, you'll need to manually seed the database later.

7. Tap or click Save to create the mailbox database copy, and then click Close when the process completes. If an error occurred, you need to take the appropriate corrective action. Otherwise, you can now work with the database copy.

In the Exchange Management Shell, you can create mailbox database copies by using the Add-MailboxDatabaseCopy cmdlet. Listing 3-5 provides the syntax and usage. Use the –ReplayLagTime parameter to specify how long the Exchange Replication Service should wait before replaying log files. Use the –TruncationLagTime parameter to specify how long the Exchange Replication Service should wait before truncating logs that have been replayed.

TIP **Different database copies can have different lag times. If you want logs to be replayed immediately, set a relatively short replay lag time or none at all. If you want a cushion for protection against inadvertent changes, set a longer replay lag time. As an example, if you have three database copies, you might want two copies to have short replay lag times and one copy to have a long replay lag time.**

NOTE **The new database copy will remain in a Suspended state if you use the –SeedingPostponed parameter. When the database copy status is set to Suspended, the SuspendMessage is set to "Replication is suspended for database copy '<Name>' because database needs to be seeded." You can seed the database as discussed in the "Updating Mailbox Database Copies" section.**

LISTING 3-5 Add-MailboxDatabaseCopy cmdlet syntax and usage

Syntax

```
Add-MailboxDatabaseCopy -Identity SourceDatabase
-MailboxServer TargetServer [-ActivationPreference PrefValue]
[-ReplayLagTime Days.Hours:Minutes:Seconds]
[-SeedingPostponed <$true | $false>]
[-TruncationLagTime Days.Hours:Minutes:Seconds]
[-DomainController FullyQualifiedName]
```

Usage

```
Add-MailboxDatabaseCopy -Identity "Engineering Primary Database"
-MailboxServer "MailServer36" -ReplayLagTime 00:03:00
-TruncationLagTime 00:10:00 -ActivationPreference 2
```

Setting replay, truncation, and preference values for database copies

Replay and truncation values are designed to let you fine-tune the way replication works for each database copy. Replay lag time is the amount of time to delay log replay. Truncation lag time is the amount of time that you want to delay log truncation after a log has been successfully replayed. You can also set a relative preference value for database copies. The preference value sets the order of activation preference after a failure or outage affecting the active database, with a value of 1 indicating the highest preference, a value of 2 the next highest preference, and so on. You cannot set a database copy to a value higher than the number of database copies.

Active Manager uses the preference value in the case of a tie during the best-copy selection process.

To set preference values for a database copy, complete the following steps:

1. In the Exchange Admin Center, select Servers in the feature pane, and then select Databases. In the main pane, you should see a list of active databases that are available in the Exchange organization.
2. Select the mailbox database with the copy that you want to manage. In the details pane, you'll see a list of active and passive copies of the database, along with separate management options for each. Tap or click the View Details option for the database copy you want to modify.
3. The current activation preference number and replay lag time are listed. Change the values as necessary, and then tap or click Save.

In the Exchange Management Shell, you can set replay, truncation, and preference values for mailbox database copies by using the Set-MailboxDatabaseCopy cmdlet. Listing 3-6 provides the syntax and usage.

LISTING 3-6 Set-MailboxDatabaseCopy cmdlet syntax and usage

Syntax

```
Set-MailboxDatabaseCopy -Identity Database\Server
[-ActivationPreference PrefValue]
[-ReplayLagTime Days.Hours:Minutes:Seconds]
[-TruncationLagTime Days.Hours:Minutes:Seconds]
[-DomainController FullyQualifiedName]
```

Usage

```
Set-MailboxDatabaseCopy -Identity "Tech Mail Database\MailServer36"
-ReplayLagTime 00:02:00 -TruncationLagTime 00:05:00
-ActivationPreference 6
```

In Exchange 2013, lagged copies automatically play down log files as necessary to accommodate adverse conditions, such as the following:

- If Exchange 2013 detects that page patching is required for a lagged copy, the logs will be automatically replayed into the lagged copy to perform page patching.
- If Exchange 2013 detects that a low disk space threshold has been reached or that no other log copies are available, the logs will be automatically replayed into the lagged copy.
- If Exchange 2013 detects that there are too few available healthy copies (active and passive) of a database for more than 24 hours, the logs will be automatically replayed into the lagged copy.

However, you must enable these options specifically. You enable lagged copy replay for all lagged copies in a particular database availability group by using the following command:

```
Set-DatabaseAvailabilityGroup DAGName -ReplayLagManagerEnabled $true
```

DAGName is the name of the database availability group to configure.

You specify the low diskspace threshold as a percentage of free disk space before log replay occurs by using the registry value:

HKLM\Software\Microsoft\ExchangeServer\v15\Replay\Parameters\ReplayLagPlayDownPercentDiskFreeSpace

For example, if you want Exchange to automatically play down lagged copies when the free disk space on the volume used by the active database reaches 10 percent, you'd edit the ReplayLagPlayDownPercentDiskFreeSpace value in the registry and set it to 10.

You specify the number of available healthy copies that triggers replay by using the following registry value:

HKLM\Software\Microsoft\ExchangeServer\v15\Replay\Parameters\ReplayLagManagerNumAvailableCopies

By default, this value is set to 3, meaning replay is triggered whenever there are fewer than 3 copies of a database available for a 24-hour period. Although this value may work well in large deployments of Exchange, this value is not ideal for small deployments. Specifically, in deployments in which you have three or fewer Mailbox servers in a DAG, setting this value to 3 will cause lagged logs to play down every 24 hours whether you want them to or not.

NOTE **As lagged copies can use SafetyNet, recovery or activation of lagged copies is much easier than in Exchange 2010. Exchange 2013 also issues single copy alerts as part of the managed availability architecture. Previously, single copy alert was implemented as a script that ran periodically as a scheduled task.**

Suspending and resuming replication

As part of planned maintenance or for other reasons, it might be necessary to temporarily suspend replication activity for a database copy. In addition, prior to performing some administrative tasks, you need to suspend replication activity before you can complete the task—for example, before performing seeding. You can suspend and resume database copy activity by completing the following steps:

1. In the Exchange Admin Center, select Servers in the feature pane, and then select Databases. In the main pane, you should see a list of active databases that are available in the Exchange organization.
2. Select the mailbox database with the copy that you want to manage. In the details pane, you'll see a list of active and passive copies of the database. Whereas the active copy of a database is normally listed with a status of Active Mounted or Active Dismounted, passive copies are normally listed with a status of Passive Healthy.
3. Select the Suspend option for the passive database copy (not an active database) for which you want to suspend replication.

4. In the Suspend Database Copy dialog box, enter a comment as to why you are suspending replication. If you want to ensure replication can only be activated by you or another administrator, select This Copy Can Be Activated Only By Manual Intervention.
5. Tap or click Save to suspend continuous replication.

To resume replication later, tap or click Resume. If a suspend comment was provided, you can read the comment. Then tap or click Yes to resume continuous replication.

In the Exchange Management Shell, you can suspend and resume replication by using Suspend-MailboxDatabaseCopy and Resume-MailboxDatabaseCopy, respectively. Listings 3-7 and 3-8 provide the syntax and usage. If you only suspend activation by using the –ActivationOnly parameter, the database cannot be activated until you resume replication without specifying the –ReplicationOnly parameter. The –ReplicationOnly parameter resumes replication without affecting the activation setting. For example, if the –ActivationSuspended parameter was set to $True, the parameter remains set to $True.

LISTING 3-7 Suspend-MailboxDatabaseCopy cmdlet syntax and usage

Syntax

```
Suspend-MailboxDatabaseCopy -Identity Database\Server
[-ActivationOnly <$true | $false>] [-EnableReplayLag <$true | $false>]
[-DomainController FullyQualifiedName]
[-SuspendComment Comment]
```

Usage

```
Suspend-MailboxDatabaseCopy -Identity "Tech Mail Database\MailServer36"
-ActivationOnly
```

LISTING 3-8 Resume-MailboxDatabaseCopy cmdlet syntax and usage

Syntax

```
Resume-MailboxDatabaseCopy -Identity Database\Server
[-DisableReplayLag <$true | $false>] [-DisableReplayLagReason "Comment"]
[-ReplicationOnly <$true | $false>] [-DomainController FullyQualifiedName]
```

Usage

```
Resume-MailboxDatabaseCopy -Identity "Tech Mail Database\MailServer36"
```

Activating lagged database copies

Lagged database copies have a replay lag time great than 0. You can activate a lagged copy by recovering the database copy from SafetyNet, by replaying all uncommitted log files, or by performing a point in time activation.

Before you activate a lagged copy, you may want to preserve the original files for the lagged copy. If so, you need to create a snapshot of the volumes containing the database copy and its log files by suspending replication of the lagged copy you want to activate, and then creating a shadow copy of these volumes as detailed in the steps that follow:

1. Suspend replication of the lagged copy you want to activate using the following command:

   ```
   Suspend-MailboxDatabaseCopy Database\Server -SuspendComment "Comment"
   -Confirm $False
   ```

   Database is the name of the lagged copy, *Server* is the name of the server hosting the lagged copy, and *Comment* is a descriptive comment, such as:

   ```
   Suspend-MailboxDatabaseCopy "Engineering DB\MailServer18"
   -SuspendComment "Suspending replication to take a db snapshot"
   -Confirm $False
   ```

2. Create a snapshot of the database and log folders by using the following command:

   ```
   Vssadmin create shadow /For="c:\Databases\Engineering DB"
   Vssadmin create shadow /For="c:\Logs\Engineering DB"
   ```

3. Optionally, copy the database and log files to another volume where you want to perform the recovery.

To recover a lagged copy from SafetyNet, complete the following steps:

1. Because you don't want the log files to replay when the database is mounted, move the log files for the database copy to an archive folder. This preserves the log files in case they are subsequently needed.
2. To allow the database to mount without all the necessary transaction logs files, you'll need to confirm that you accept the data loss. To do this, mount the database with the -AcceptDataLoss parameter as shown in this example:

   ```
   Mount-Database "Engineering DB" -AcceptDataLoss
   ```

3. Exchange will mount the database and then request redelivery of missing messages from SafetyNet. You can confirm that the lagged copy was successfully activated by viewing the database properties. In Exchange Admin Center, select Servers in the feature pane, and then select Databases. Next, select the database copy you activated. In the Details pane, click View Details.
4. After you verify that the database copy was successfully activated, you can delete the log files you moved to an archive folder, as these logs are no longer needed.

To activate a lagged copy by replaying all uncommitted log files, complete the following steps:

1. Activate the lagged copy on a specified server by using the following command:

   ```
   Move-ActiveMailboxDatabase Database -ActivateOnServer Server
   -SkipLagChecks
   ```

   Database is the name of the lagged copy and *Server* is the name of the server hosting the lagged copy, such as:

   ```
   Move-ActiveMailboxDatabase "Engineering DB" -ActivateOnServer
   MailServer18 -SkipLagChecks
   ```

2. Exchange will mount the database on the designated server and replay all the log files. The duration of the replay process depends on the amount of data to replay and the speed at which your server hardware can replay the logs.
3. You can confirm that the lagged copy was successfully activated by viewing the database properties. In Exchange Admin Center, select Servers in the feature pane, and then select Databases. Next, select the database copy you activated. In the Details pane, click View Details.

To activate a lagged copy to a point in time, complete the following steps:

1. Before you can activate a lagged copy to a point in time, you must first determine which log files are required to meet your recovery requirements. Use the log file date and time to identify which log files you need and which log files should be moved to an archive directory until the recovery process is successfully completed. Specifically, any log file created after your recovery time should be moved to the archive directory.
2. Next, you need to delete the checkpoint file for the lagged copy. This file has the .chk extension.
3. At an elevated command prompt, use Eseutil to perform the recovery operation. The basic syntax is:

   ```
   Eseutil /r ENN /a
   ```

   ENN is the log generation prefix for the database, such as E00 or E01. This prefix is used with all the database files, so it's easily identified when you access the database folder for the lagged copy.
4. When all the logs have been replayed, the database will be in a clean state and you can optionally copy the database and log files to another volume where you want to perform the recovery. Keep in mind that the duration of the replay process depends on the amount of data to replay and the speed at which your server hardware can replay the logs.
5. You can confirm that the lagged copy was successfully activated by viewing the database properties. In Exchange Admin Center, select Servers in the feature pane, and then select Databases. Next, select the database copy you activated. In the Details pane, click View Details.

Updating mailbox database copies

Seeding is the process of initially replicating an active or passive database into a database copy. This creates a baseline passive copy of a database. Normally, seeding occurs automatically, and the length of time required to completely seed a database depends on the size of the source database, the available bandwidth on the network, and the level of activity on the servers involved. However, automatic seeding can fail, and in this case, you then need to manually initiate seeding.

> ***REAL WORLD*** **An automatic seed produces a copy of an active or passive database on a target Mailbox server. Automatic seeding occurs only during the creation of a new database or for a database that has never been backed up.**
>
> **You can identify a problem with seeding by checking the state of the database copy. When you create a database copy, the database should enter the Initializing state and then the Seeding state. When seeding is complete, the database copy should be in the Healthy state. If the database remains in a Suspended state and does not complete initialization or seeding, there is a problem. Note also that if you are seeding when creating the copy, the task will not complete successfully until the seed is completed. So, you simply watch the task progress and do not need to check copy status.**

You can reseed a mailbox database copy anytime you suspect divergence has occurred. However, divergence isn't necessarily a problem because incremental reseed (incremental resync) takes care of resolving the divergence. You would not need to do a full reseed except in circumstances in which resync isn't possible—for example, when there is no overlap in log files between diverged copies, or when you've done something you shouldn't have, like an offline defragmentation of a copy that causes uncorrectable divergence.

When you reseed a database, Exchange empties the database copy and replicates a new passive database copy. Typically, you won't need to reseed database copies after the initial seeding has occurred; however, in some situations you might need to reseed a database copy. One state you can check for is the FailedAnd Suspended state. In this state, Exchange has detected a failure and suspended replication replay because resolution of the failure explicitly requires administrator intervention. For example, if Exchange detects an unrecoverable divergence between the active mailbox database and a database copy, Exchange marks the database copy as FailedAndSuspended. If an incremental resync doesn't eventually resolve the problem, you need to resolve the underlying cause of the failure before the database copy can be transitioned to a healthy state, which includes reseeding the database.

Before you can seed or reseed a database, you must suspend replication. For very large databases—that is, those that are multiple terabytes (TB) in size—the preferred technique for seeding the initial passive copy of the database, if service level agreements allow or such an outage is acceptable, is to dismount the active copy of the database and copy the database file to the same location on the target Mailbox server in the same database availability group. Rather than copying the database

over the network, which could take several days for a multiterabyte database, you should consider the following:

- Copying the database to one or more disk drives, preferably hot-swappable drives that can be moved between the source and target servers
- Copying the database to one or more logical unit numbers (LUNs) in your storage array that can be assigned to or is assigned to the target server

With this approach, the database will be unavailable until seeding is completed and you can mount the database. Alternatively, you can leave the active database online and use the Exchange Management tools to initiate the seeding process. After you've created at least one baseline passive copy of a database, you can seed new passive copies from the baseline passive copy at any time by using an online or offline approach.

The size of the database, the available network bandwidth, network latency, and the activity levels on the source and target servers determine how long an over-the-network transfer or update takes. After the seeding process has started, don't close the Exchange Admin Center or Exchange Management Shell until the process has completed. If you do, the seeding operation will be terminated and will need to be restarted.

Keep the following in mind when you are considering updating database copies:

- When you seed a database using the Exchange Admin Center, both the database copy and the content index catalog are seeded. In the Exchange Management Shell, you can specify that only the database copy should be seeded using the –DatabaseOnly parameter or that only the context index catalog should be seeded using the –CatalogOnly parameter.
- Before you seed the database copy, you may want to manually remove existing files on the server that hosts the database copy. You can delete existing files in the Exchange Management Shell by using the –DeleteExistingFiles parameter; however, these options remove only the files Exchange checks for and might fail if other files are present.
- When seeding is complete, Exchange automatically resumes replication. If you want to resume replication manually instead, you can use the –ManualResume parameter in the Exchange Management Shell.
- By default, seeding data is transferred over the replication network for the database availability group, unless you are seeding to a remote site, in which case it will default to the messaging network. You can override the defaults by using the –Network parameter. The network compression and encryption settings are used and determine whether the transferred data is compressed, encrypted, or both. You specify the networks to use by name in both management tools. In the Exchange Management Shell, you can override the network compression and encryption settings by using –NetworkCompression-Override and –NetworkEncryptionOverride, respectively.

You can seed a database manually by completing the following steps:

1. In the Exchange Admin Center, select Servers in the feature pane, and then select Databases. In the main pane, you should see a list of active databases that are available in the Exchange organization.
2. Select the mailbox database with the copy that you want to manage. In the details pane, you'll see a list of active and passive copies of the database.
3. Select the Update option for the passive database copy (not an active database) that you want to update (see Figure 3-8).

TIP **The Exchange Admin Center won't let you reseed a database that's in a healthy or other normal state. However, you can force a reseed by suspending the database, copying it, and then updating the database copy.**

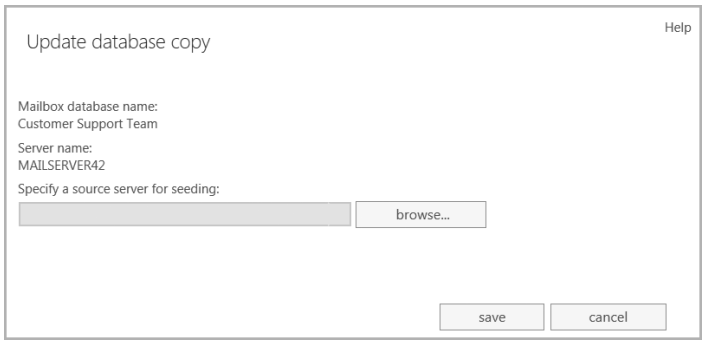

FIGURE 3-8 Update a database copy.

4. By default, Exchange will seed the database from the active copy of the database. If you want to use a passive copy for seeding, tap or click Browse. In the Select Mailbox Server dialog box, select the source server hosting the passive copy you want to use and then tap or click OK.
5. Tap or click Save to begin seeding. If an error occurred, you need to take the appropriate corrective action. Tap or click Close.

To seed a database copy in the Exchange Management Shell, you use the Update-MailboxDatabaseCopy cmdlet. Listing 3-9 provides the syntax and usage. Use the –Force parameter when seeding programmatically, and you will not be prompted for administrative input.

LISTING 3-9 Update-MailboxDatabaseCopy cmdlet syntax and usage

Syntax

```
Update-MailboxDatabaseCopy -Identity Database\Server
-SourceServer ServerName [-CancelSeed <$true | $false>]
[-BeginSeed <$true | $false>] [-CatalogOnly <$true | $false>]
[-DatabaseOnly <$true | $false>] [-DeleteExistingFiles <$true | $false>]
[-DomainController FullyQualifiedName] [-Force <$true | $false>]
[-ManualResume <$true | $false>] [-MaximumSeedsInParallel MaxNumSeeds]
[-NetworkCompressionOverride {"UseDAGDefault"|"Off"|"On"}]
[-NetworkEncryptionOverride {"UseDAGDefault"|"Off"|"On"}]
[-Network NetworkID] [-SafeDeleteExistingFiles <$true | $false>]
```

Usage

```
Update-MailboxDatabaseCopy -Identity "CS Mail\MailServer25"
-CatalogOnly -Force

Update-MailboxDatabaseCopy -Identity "CS Mail\MailServer25"
-DatabaseOnly

Update-MailboxDatabaseCopy -Identity "CS Mail\MailServer25"
-Network "EastCampusDAG1\Primary DAG Network"
-NetworkCompressionOverride "On" -NetworkEncryptionOverride "Off"
```

Monitoring database replication status

As an Exchange administrator, you need to monitor the health and status of database copies to ensure that they are available when needed. You can view key health and status information for a database copy by completing the following steps:

1. In the Exchange Admin Center, select Servers in the feature pane, and then select Databases. In the main pane, you should see a list of active databases that are available in the Exchange organization.
2. Select the mailbox database with the copy that you want to manage. In the details pane, you'll see a list of active and passive copies of the database along with the current status of each. Table 3-1 lists the possible status values and any corrective action that might be required.
3. Tap or click the View Details option for the database copy with which you want to work. This opens a properties dialog box.
4. Use the information provided to determine the health and status of replication for the database copy. The information provided includes:
   - **Database** Displays the name of the selected database
   - **Mailbox Server** Displays the name of the Mailbox server that hosts the database copy
   - **Content Index State** Displays the status of content indexing for the database copy

- **Status** Displays the current health and status of replication for the database copy
- **Copy Queue Length** Shows the number of log files waiting to be copied and checked
- **Replay Queue Length** Shows the number of log files waiting to be replayed into this copy of the database
- **Error Messages** Displays any current error status or error message for the database copy
- **Latest Available Log Time** Shows the time associated with the latest available log generated by the active database copy
- **Last Inspected Log Time** Shows the modification time of the last log that was successfully validated by the Mailbox server hosting the database copy
- **Last Copied Log Time** Shows the modification time of the last log that was successfully copied
- **Last Replayed Log Time** Shows the modification time of the last log that was successfully replayed by the Mailbox server hosting the database copy
- **Activation Preference Number** Shows the activation preference value for the database copy
- **Replay Lag Time** Shows the current replay lag time in days, if any

In the Exchange Management Shell, you can check the health and status of replication by using the Get-MailboxDatabaseCopyStatus cmdlet. Listing 3-10 provides the syntax and usage.

LISTING 3-10 Get-MailboxDatabaseCopyStatus cmdlet syntax and usage

Syntax

```
Get-MailboxDatabaseCopyStatus -Server ServerName {AddtlParams}

Get-MailboxDatabaseCopyStatus [-Identity LocalDatabaseName]
[-Active <$true | $false>] [-Local <$true | $false>]] {AddtlParams}

{AddtlParams}
[-ConnectionStatus <$true | $false>] [-DomainController FullyQualifiedName]
[-ExtendedErrorInfo <$true | $false>] [-UseServerCache <$true | $false>]
```

Usage

```
Get-MailboxDatabaseCopyStatus –Server "MailServer35"
-ConnectionStatus –ExtendedErrorInfo

Get-MailboxDatabaseCopyStatus

Get-MailboxDatabaseCopyStatus –Identity "Accounting Mail"
```

TABLE 3-1 Copy status values

COPY STATUS	STATUS OF THE MAILBOX DATABASE COPY	CORRECTIVE ACTION
Activation Suspended	Has been manually blocked from activation by an administrator.	Allow activation, if appropriate.
Disconnected And Healthy	Has been disconnected, and was in the Healthy state when the loss of connection occurred.	This state can be reported during network failures between the active copy and the database copy.
Disconnected And Resynchronizing	Is no longer connected to the active database copy, and was in the Resynchronizing state when the loss of connection occurred.	This state can be reported during network failures between the active copy and the database copy.
Dismounted	Is offline and not accepting client connections. Applies only to the active mailbox database.	Mount the database if maintenance is complete.
Dismounting	Is going offline and terminating client connections. Applies only to the active mailbox database.	N/A
Failed	Is in a failed state because it is not suspended, and is not able to copy or replay log files.	Exchange periodically checks to see whether the problem that caused the copy status to change to Failed has been resolved. If so, and barring no other issues, the copy status automatically changes to Healthy.
Failed And Suspended	Is in the Failed And Suspended state because a failure was detected and because resolution of the failure explicitly requires administrator intervention.	Take corrective action as appropriate. Exchange does not periodically check to see whether the problem has been resolved and does not automatically recover.

COPY STATUS	STATUS OF THE MAILBOX DATABASE COPY	CORRECTIVE ACTION
Healthy	Is successfully copying and replaying log files, or has successfully copied and replayed all available log files.	N/A
Initializing	Is being created, or the Microsoft Exchange Replication service is starting up or has just been started, or the Mailbox Database copy is transitioning to another state.	While the copy is in this state, Exchange is verifying that the database and log stream are in a consistent state. It should generally not be in this state for longer than 30 seconds.
Mounted	Is online and accepting client connections. Applies only to active mailbox database.	N/A
Mounting	Is coming online and not yet accepting client connections. Applies only to active mailbox database.	N/A
Resynchronizing	Is being checked for any divergence between the active copy and this passive copy.	The copy status remains in this state until any divergence is detected and resolved.
Seeding	Is being seeded, the related content index is being seeded, or both.	Upon successful completion of seeding, the copy status should change to Initializing.
Service Down	Cannot connect to the replication service.	Start or restart the Microsoft Exchange Replication service on the server that hosts the mailbox database copy.
Single Page Restore	Had a single page error, and this error is being corrected automatically.	N/A
Suspended	Is in a suspended state as a result of an administrator manually suspending the database copy.	Resume replication if appropriate

Removing database copies

You can remove a passive database copy at any time by using the Exchange Management tools. After removing a database copy, you need to manually delete any database and transaction log files from the server.

> ***NOTE*** **You cannot use these procedures to remove the active copy of a mailbox database. To remove a database that is an active copy, you must first switch the database over to a new active copy. Alternatively, if you no longer want a database and its copies, you first need to remove all passive copies, and then you need to remove all mailboxes from the active database before you can delete it.**

> ***TIP*** **You can remove mailbox database copies only from a database availability group with a Healthy status. If the database availability group doesn't have a Healthy status, you won't be able to remove any mailbox database copies.**

To remove a database copy, complete the following steps:

1. In the Exchange Admin Center, select Servers in the feature pane, and then select Databases. In the main pane, you should see a list of active databases that are available in the Exchange organization.
2. Select the mailbox database with the copy that you want to manage. In the details pane, you'll see a list of active and passive copies of the database along with the current status of each.
3. Tap or click the Remove option for the database copy you want to remove.
4. When prompted to confirm, tap or click Yes. If an error occurred, you need to take the appropriate corrective action. Tap or click Close.

In the Exchange Management Shell, you can remove a database copy by using the Remove-MailboxDatabaseCopy cmdlet. Listing 3-11 provides the syntax and usage.

LISTING 3-11 Remove-MailboxDatabaseCopy cmdlet syntax and usage

Syntax

```
Remove-MailboxDatabaseCopy -Identity DatabaseName\ServerName]
[-DomainController FullyQualifiedName]
```

Usage

```
Remove-MailboxDatabaseCopy –Identity "CS Mail Database\MailServer24"
```

Before you remove a database using the shell, you may need to identify available copies of the database. To do this, enter the following command:

```
Get-MailboxDatabase DatabaseName | fl DatabaseCopies
```

Where *DatabaseName* is the name of the database you want to work with, such as:

```
Get-MailboxDatabase Development | fl databasecopies
```

As shown in this example, the output lists where copies of the database are available:

```
DatabaseCopies : {Development\MAILSERVER42, Development\MAILSERVER21}
```

Managing mailbox databases

Now that you know how to create and use databases, let's look at some general techniques you'll use to manage databases.

> ***NOTE*** **These techniques apply only to active mailbox databases. Passive copies of mailbox databases are managed as discussed in the "Working with mailbox database copies" section earlier in this chapter.**

Mounting and dismounting databases

You can only access databases that are mounted. If a database isn't mounted, the database isn't available for use. If a database isn't mounted it means that an administrator has probably dismounted the database or that the drive on which the database is located isn't online. It could also mean that the Exchange Information Store service is not running or that the drive, log drive, or both are online but out of disk space.

> ***REAL WORLD*** **A dismounted database can also indicate that there are problems with the database, transaction log, and system files used by the database. During startup, Exchange Server 2013 obtains a list of database files registered in Active Directory and then checks for the related files before mounting each database. If files are missing or corrupted, Exchange Server 2013 will be unable to mount the database. Exchange Server 2013 then generates an error and logs it in the application event log on the Exchange server. A common error is Event ID 9547. An example of this error follows:**
>
> ```
> The Active Directory indicates that the database file
> D:\Exchsrvr\mdbdata\Marketing.edb exists for the Microsoft
> Exchange
> Database; however, no such files exist on the disk.
> ```
>
> **This error tells you that the Exchange database (Marketing.edb) is registered in Active Directory but Exchange Server 2013 is unable to find the file on the disk. When Exchange Server 2013 attempts to start the corrupted mailbox database, you'll see an additional error as well. The most common error is Event ID 9519. An example of this error follows:**
>
> ```
> Error 0xfffffb4d starting database Marketing on the Microsoft
> Exchange Information Store.
> ```
>
> **This error tells you that Exchange Server 2013 couldn't start the Marketing database. You can try to restore the database to recover its data. If you are unable to restore the database file, you can create a copy of all database files and store them elsewhere and then recreate the database structures in the Exchange Admin Center by mounting the database. When you mount the database, Exchange Server 2013 creates a new database file. As a result, the data in the original database files (and not the copies) is lost and cannot be recovered. Exchange Server 2013 displays a warning before mounting the database and recreating the database file. Tap or click Yes only when you are absolutely certain that you cannot recover the database.**

Be sure you don't overwrite the database files containing the data you want to try to recover. You can still work on the database while users access the newly created empty database. This is effectively a dial-tone database that you are creating. Then, take the damaged database file elsewhere, run repair, make the database consistent, and then use it to complete the dial-tone recovery process.

If you can't restore or repair a database and you need as much of the data as you can get back, you might have clients in cached or offline mode with viable copies of the data that can be exported and imported.

Determining the status of databases

Mailbox databases have several associated states, including the following:

- Mounted
- Backup In Progress
- Background Database Maintenance

You can determine the status of a database by following these steps:

1. In the Exchange Admin Center, select Servers in the feature pane, and then select Databases. In the main pane, you should see a list of active databases that are available in the Exchange organization.
2. Select the mailbox database that you want to work with to see a list of all copies of that database in the details pane. The status of each database copy also is listed in the details pane.

In the Exchange Management Shell, you can determine the status of all databases or specific databases by using the Get-MailboxDatabase. Listing 3-12 provides the syntax and usage for this cmdlet. To see status details, you can specify the status flags associated with each state you want to see as part of the formatted output. In the example, the Mounted, Backup In Progress, and Background Database Maintenance status values are listed as True or False.

LISTING 3-12 Getting database status details

Syntax

```
Get-MailboxDatabase [-Identity MailboxDatabase | -Server Server]
[-DomainController DCName] [-DumpsterStatistics <$true | $false>]
[-IncludePreExchange2013 <$true | $false>]
[-Status <$true | $false>] | format-table Name, Mounted, BackupInProgress,
BackgroundDatabaseMaintenance
```

Usage for specific database

```
Get-MailboxDatabase –Identity "Eng DB" –Status | format-table Name,
Mounted, BackupInProgress, BackgroundDatabaseMaintenance
```

Usage for all databases on a server

```
Get-MailboxDatabase –Server "CORPSVR127" –Status | format-table
Name, Mounted, BackupInProgress, BackgroundDatabaseMaintenance
```

Usage for all databases

```
Get-MailboxDatabase –Status | format-table Name,
Mounted, BackupInProgress, BackgroundDatabaseMaintenance
```

Dismounting and mounting databases

Before you perform maintenance on a Mailbox server in a database availability group, you should perform a server switchover so that the server's active databases are transitioned and made active on one or more additional servers in the group. You might also want to suspend replication or block activation of passive copies on the server being maintained. For mailbox databases that are not part of an availability group, you should rarely dismount an active database, but if you need to do so, follow these steps:

1. In the Exchange Admin Center, select Servers in the feature pane, and then select Databases. In the main pane, you should see a list of active databases that are available in the Exchange organization.
2. Select the mailbox database that you want to copy to see a list of all copies of that database in the details pane. Note the status of the active database copy. If the copy is mounted, the status normally is listed as Active Mounted.
3. Tap or click the More button (this button shows three dots), and then select Dismount. When prompted, confirm the action by tapping or clicking Yes. Exchange Server dismounts the database. Users will no longer be able to access the database and work with their server-based folders.

After you've dismounted a database and performed maintenance, recovery, or other procedures as necessary, you can remount the database by selecting the database, tapping or clicking the More button, and then selecting Mount. When prompted, confirm the action by tapping or clicking Yes.

In the Exchange Management Shell, you can dismount and mount databases by using the Dismount-Database and Mount-Database cmdlets, respectively. Listing 3-13 provides the syntax and usage for these cmdlets.

LISTING 3-13 Dismounting and mounting databases

Syntax

```
Dismount-Database –Identity DatabaseIdentity
[-DomainController FullyQualifiedName]

Mount-Database –Identity DatabaseIdentity
[-AcceptDataLoss <$true | $false>] [-DomainController FullyQualifiedName]
[-Force <$true | $false>]
```

Usage for dismounting a database

```
Dismount-Database –Identity "Eng DB"
```

Usage for mounting a database

```
Mount-Database –Identity "Eng DB"
```

Specifying whether a database should be automatically mounted

Normally, Exchange Server automatically mounts databases on startup. You can, however, change this behavior. For example, if you're recovering an Exchange server from a complete failure, you might not want to mount databases until you've completed recovery. In this case, you can disable automatic mounting of databases.

To enable or disable automatic mounting of a database, complete the following steps:

1. In the Exchange Admin Center, select Servers in the feature pane, and then select Databases. In the main pane, you should see a list of active databases that are available in the Exchange organization.
2. Double-tap or double-click the database with which you want to work.
3. On the Maintenance page, do one of the following and then tap or click Save:
   a. To ensure that a database isn't mounted on startup, select the Don't Mount This Database At Startup check box.
   b. To mount the database on startup, clear the Don't Mount This Database At Startup check box.

In the Exchange Management Shell, you can enable or disable automatic mounting at startup by using the Set-MailboxDatabase. Listing 3-14 provides the syntax and usage for controlling automatic mounting.

LISTING 3-14 Controlling automatic mounting

Syntax

```
Set-MailboxDatabase –Identity DatabaseIdentity
 –MountAtStartup <$true | $false>
```

Usage

```
Set-MailboxDatabase –Identity "Eng DB"
 –MountAtStartup $false
```

Setting the maintenance interval

You should run maintenance routines against databases on a daily basis. The maintenance routines organize the databases, clear out extra space, and perform other essential housekeeping tasks. By default, the automatic background maintenance does some of this work, and Exchange Server runs extended, foreground maintenance tasks daily from 1:00 A.M. to 5:00 A.M. If this conflicts with other scheduled administrative tasks or activities on the Exchange server, you can change the maintenance settings by following these steps:

1. In the Exchange Admin Center, select Servers in the feature pane, and then select Databases. In the main pane, you should see a list of active databases that are available in the Exchange organization.
2. Double-tap or double-click the database with which you want to work. This opens a properties dialog box for the database.
3. On the Maintenance page, note the current Maintenance Schedule, and then select Customize. You can now set the times when maintenance should occur by using the options in the Customize Maintenance dialog box, shown in Figure 3-9.
   - Times that are used for maintenance are filled in with a dark bar.
   - Times that aren't used for maintenance are blank.

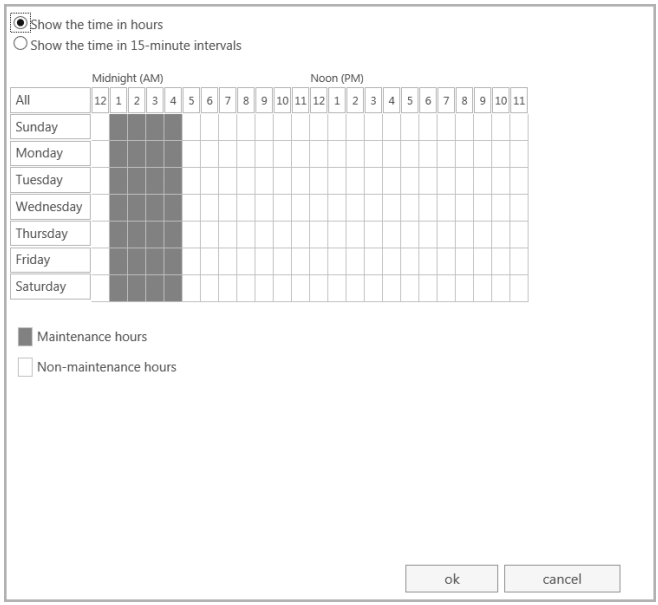

FIGURE 3-9 Customize the background maintenance schedule.

IMPORTANT **Ideally, you want to schedule background maintenance to occur during off-peak hours. As the size of databases and activity levels change, you'll want to optimize this schedule to ensure that servers aren't overburdened and that servers have enough time to perform necessary maintenance tasks.**

4. Show the time in hours or in 15-minute intervals using the options provided. Change the setting for a time interval by tapping or clicking it.
   - Hourly or 15-minute interval buttons are used to select or clear a particular interval for all the days of a week.
   - Days of the week buttons allow you to clear or select all the hours in a particular day.
   - The All button allows you to clear or select all the time periods.
5. Tap or click OK to close the Customize Maintenance Schedule dialog box.
6. By default, Exchange performs background maintenance tasks by scanning the ESE 24 hours a day, seven days a week. Select or clear the related check box as appropriate. Note that if you change this setting, you must dismount and then remount the database for the change to take effect. Tap or click OK.

In the Exchange Management Shell, you can configure the maintenance schedule for a database by using Set-MailboxDatabase. Listing 3-15 provides the syntax and usage. In the example, replication is configured to occur between Friday at 9:00 P.M. and Monday at 1:00 A.M.

LISTING 3-15 Setting the maintenance schedule

Syntax

```
Set-MailboxDatabase –Identity DatabaseIdentity
[–MaintenanceSchedule Schedule]
[-BackgroundDatabaseMaintenance <$true | $false>]
```

Usage

```
Set-MailboxDatabase –Identity "Eng DB"
 –MaintenanceSchedule "Fri.9:00 PM-Mon.1:00 AM"
```

Moving databases

As discussed earlier, each database has a database file associated with it, and the location of this file has an important role in managing Exchange Server performance. You can view the current file and folder paths the database is using by entering:

```
Get-MailboxDatabase DatabaseName | fl *path
```

DatabaseName is the name of the database to check, such as:

```
Get-MailboxDatabase "Engineering" | fl *path
```

In the command output, the database file path is listed as EdbFilePath, the log folder path is listed as LogFolderPath and any associated temporary data folder is listed under TemporaryDataFolderPath, as shown in this example:

```
EdbFilePath             : G:\Databases\Engineering\Engineering.edb
LogFolderPath           : H:\Logs\Engineering
TemporaryDataFolderPath :
```

In the Exchange Management Shell, you can move databases by using the Move-DatabasePath cmdlet. Listing 3-16 provides the syntax and usage. If the specified database is mounted, the database is automatically dismounted and then remounted, and it is unavailable to users while it's dismounted. Additionally, you can perform a database move only while logged on to the affected Mailbox server, with one exception. If you are performing a configuration-only move, you can perform the configuration-only move from your management computer.

LISTING 3-16 Move-DatabasePath cmdlet syntax and usage

Syntax

```
Move-DatabasePath -Identity DatabaseIdentity
[-ConfigurationOnly <$true | $false>] [-EdbFilePath EdbFilePath]
[-DomainController DCName] [-Force <$true | $false>]
[-LogFolderPath FolderPath]
```

Usage

```
Move-DatabasePath -Identity "Engineering"
-EdbFilePath "K:\Databases\Engineering\Engineering.edb"
-LogFolderPath "L:\Logs\Engineering"
```

You cannot move a database that is being backed up or a replicated mailbox database. To move a replicated mailbox database, you must disable circular logging if enabled, remove all replicated copies, and then perform the move operation. After the move is complete, you can add copies of the mailbox database and re-enable circular logging. You'll also want to rebuild the content indexes for each copy of the database. To perform these and other related tasks, complete the following steps:

1. Identify any replay lag or truncation lag settings for all copies of the mailbox database being moved by entering the following command:

   ```
   Get-MailboxDatabase DatabaseName | fl *lag*
   ```

   DatabaseName is the name of the database that you want to move.
2. Disable circular logging if the option is enabled by entering the following command:

   ```
   Set-MailboxDatabase DatabaseName -CircularLoggingEnabled $false
   ```

3. Identify all copies of the database by entering the following command:

   ```
   Get-MailboxDatabase DatabaseName | fl DatabaseCopies
   ```

4. Remove the mailbox database copies by entering the following command for each copy:

   ```
   Remove-MailboxDatabaseCopy DatabaseName\ServerName
   ```

   DatabaseName is the name of the database copy to remove and *ServerName* is the name of the server.
5. On each server that hosted a copy of the database, move the data files and log files for the database copy to a local archive folder, such as C:\Archives\Database for the database files and C:\Archives\Logs for the log files. This preserves the files on the server so that the database copies don't need to be reseeded after they have been recreated.
6. Use the Move-DatabasePath cmdlet to move the database path and log path to a new location. The syntax is:

   ```
   Move-DatabasePath -Identity DatabaseName
   -EdbFilePath EdbFilePath -LogFolderPath FolderPath
   ```

   During the move operation, the database will be dismounted. When Exchange finishes moving the database, Exchange will automatically mount the database.
7. On each server that hosted a passive copy of the database, create the required folders for the database and logs. For example, if you moved the database to K:\Databases\Engineering\Engineering.edb, you must create the K:\Databases\Engineering folder on each server. If you moved the log folder to L:\Logs\Engineering, you must also create the L:\Logs\Engineering folder on each server. Because the active copy of the database was moved already, you don't need to create folders for the active copy.

8. On each server that hosted a passive copy of the database, move the preserved database files to the database folder and then move the preserved log files to the log folder. As the active copy of the database was moved already, you don't need to move the preserved files for the active database.
9. Use the Add-MailboxDatabaseCopy cmdlet to add a passive copy of the database to each server that previously hosted a passive copy of the database. The basic syntax is:

   ```
   Add-MailboxDatabaseCopy -Identity SourceDatabase
   -MailboxServer TargetServer
   ```

   Don't set any replay lag or truncation lag times yet because you want to ensure the databases are recovered using the local files (and without reseeding).
10. Recreate the context indexes on each server that hosts an active or passive copy of the database. To do this, use the following commands to stop and then start the Microsoft Exchange Search service:

    ```
    Stop-Service MSExchangeSearch
    Start-Service MSExchangeSearch
    ```

11. If you want to enable circular logging of the active copy of the database, enter the following command:

    ```
    Set-MailboxDatabase DatabaseName -CircularLoggingEnabled $true
    ```

12. Use the Set-MailboxDatabseCopy cmdlet to reconfigure replay lag and truncation lag times, as appropriate. The basic syntax is:

    ```
    Set-MailboxDatabaseCopy -Identity Database\Server
    [-ReplayLagTime Days.Hours:Minutes:Seconds]
    [-TruncationLagTime Days.Hours:Minutes:Seconds]
    ```

After you've completed all these tasks, you should confirm that replication is working as expected by using the Get-MailboxDatabaseCopyStatus cmdlet. You also should use the Test-ReplicationHealth cmdlet to verify the health and status of the database availability group.

Renaming databases

To rename a database, follow these steps:

1. In the Exchange Admin Center, double-tap or double-click the database with which you want to work. This opens a properties dialog box for the database.
2. In the Name text box, type the new name for the database. Tap or click Save.

NOTE **All objects in Active Directory are located by a unique identifier. This identifier uses the directory namespace and works through each element in the directory hierarchy to a particular object. When you change the name of a database, you change the namespace for all the objects in the database.**

In the Exchange Management Shell, you can rename databases by using the –Name parameter of the Set-MailboxDatabase. Listing 3-17 provides the syntax and usage.

LISTING 3-17 Renaming a database

Syntax

```
Set-MailboxDatabase –Identity DatabaseIdentity
 –Name NewName
```

Usage

```
Set-MailboxDatabase –Identity "Eng DB"
 –Name "Engineering Mail Database"
```

Deleting databases

Before deleting a mailbox database, you must delete or move the mailboxes it contains. After you've moved items that you might need and deleted items you don't need, you can delete the database by completing the following steps:

1. In the Exchange Admin Center, select the database you want to delete, and then select the Delete button.
2. When prompted, confirm the action by tapping or clicking Yes.
3. After removing the database, you need to delete any database and transaction log files from the server.

In the Exchange Management Shell, you can delete databases by using the Remove-MailboxDatabase. Listing 3-18 provides the syntax and usage.

LISTING 3-18 Removing databases

Syntax

```
Remove-MailboxDatabase –Identity DatabaseIdentity
[-DomainController FullyQualifiedName]
```

Usage

```
Remove-MailboxDatabase –Identity "Eng DB"
```

Content indexing

Content indexing is a built-in Exchange feature. Every Exchange server in your organization supports and uses some type of indexing. To manage indexing more effectively, use the techniques discussed in this section.

Understanding indexing

Content indexing enables fast searches and lookups through server-stored mailboxes. Exchange Server 2013 uses the content indexing engine from the Microsoft Search Foundation. The Exchange Server storage engine automatically implements and manages Exchange Search. Exchange Search is used with searches for common

key fields, such as message subjects. Users take advantage of Exchange Search every time they use the Find feature in Microsoft Office Outlook. With server-based mail folders, Exchange Search is used to quickly search To, From, Cc, and Subject fields. With public folders, Exchange Search is used to quickly search From and Subject fields.

As you probably know, users can perform advanced searches in Office Outlook as well. For example, in Outlook 2010, all users need to do is tap or click in the Search box or press Ctrl+E to access the Search tools, tap or click Search Options, and then tap or click Advanced Find. In the Advanced Find dialog box, users can enter their search parameters and then tap or click Find Now.

When Exchange Server receives an advanced query for personal folders, it searches through every message in every folder. This means that as Exchange mailboxes grow, so does the time it takes to complete an advanced search. With standard searching, Exchange Server is unable to search through message attachments.

With server-based folders, Exchange Server builds an index of all searchable text in a particular mailbox database before users try to search. The index can then be updated or rebuilt at a predefined interval. Then, when users perform advanced searches, they can quickly find any text within a document or attachment.

A drawback of content indexing is that it can be resource-intensive. As with any database, creating and maintaining indexes requires CPU time and system memory, which can affect Exchange performance. Full-text indexes also use disk space. A newly created index uses approximately 10 to 20 percent of the total size of the Exchange database (and is directly related to what's in the database's mailboxes). This means that a 1-TB database would have an index of about 100 to 200 GB.

Each time you update an index, the file space that the index uses increases. Don't worry—only changes in the database are stored in the index updates. This means that the additional disk space usage is incremental. For example, if the original 1-TB database grew by 1 GB, the index could use up to 201 GB of disk space (up to 200 GB for the original index and 1 GB for the update).

Managing Exchange Store Search

Exchange Server 2013 doesn't allow administrators to configure how indexing works. Full-text indexes are stored as part of the Exchange data files. Because of this, whatever folder location you use for Exchange data files will have an indexing subfolder, which contains all the Exchange Search data for the related database and all its related databases. By default, you'll find full-text index files for a database in the %SystemDrive%\Program Files\Microsoft\Exchange Server\V15\Mailbox*Database Name**GUID* folder where GUID is the database's globally unique identifier.

NOTE **Exchange maintains full-text indexes as part of the database maintenance schedule. See the "Setting the maintenance interval" section earlier in this chapter for more information.**

Each database has an index. If you make a database copy, you are also making an index copy. There's often no need to rebuild an index. That said, as part of the recovery process for a mailbox database, you might want to rebuild the related full-text index catalog to ensure it's current. You might also want to rebuild the full-text index after you've made substantial changes to a database or if you suspect the full-text index is corrupted.

You can rebuild an index manually at any time. Exchange Server rebuilds an index by recreating it. This means that Exchange Server takes a new snapshot of the database and uses this snapshot to build the index from scratch. To manually rebuild an index, enter the following commands to stop and then start the Exchange Search service:

```
Stop-Service MSExchangeFastSearch
Start-Service MSExchangeFastSearch
```

Exchange Discovery relies on Exchange Search for databases and mailboxes within databases. You can enable or disable indexing for individual databases by setting the -IndexEnabled parameter of the Set-MailboxDatabase cmdlet to $true or $false, respectively. The following example disables indexing of the Engineering database:

```
Set-MailboxDatabase "Engineering Database" -IndexEnabled $false
```

When you disable indexing of a database, you also prevent the Exchange 2013 Discovery feature from returning messages from the database or server.

You can disable indexing for all databases on a server by stopping and disabling the Microsoft Exchange Search service. Here's an example using the Exchange Management Shell in which you stop and disable the Exchange Search service on a remote server named Server18:

```
Stop-Service MSExchangeFastSearch -ComputerName Server22

Set-Service MSExchangeFastSearch -StartupType Disabled -ComputerName
Server22
```

You can enable indexing for all databases on a server by enabling the Microsoft Exchange Search service for automatic startup and starting the service. An example using the Exchange Management Shell follows:

```
Set-Service MSExchangeFastSearch -StartupType Automatic
-ComputerName Server18

Start-Service MSExchangeFastSearch -ComputerName MailServer11
```

When you disable indexing on a server, you also prevent Exchange Discovery for all databases on the server.

Troubleshooting indexing

You can quickly determine which databases have indexing enabled by using the following command:

```
Get-MailboxDatabase | ft Name,IndexEnabled
```

You can determine whether content indexing has a healthy status by using the following command:

```
Get-MailboxDatabaseCopyStatus | ft Identity, ActiveDatabaseCopy,
ContentIndexState -Auto
```

If you find that the context index for a passive database copy is outdated, you can rebuild or reseed the index. To reseed the index, enter the following command:

```
Update-MailboxDatabaseCopy -Identity Database\Server -CatalogOnly
```

Database is the name of the database and Server is the name of the server hosting the database, such as:

```
Update-MailboxDatabaseCopy -Identity Engineering\MailServer12
-CatalogOnly
```

If you need to troubleshoot Exchange Search issues, you can use Test-Exchange Search. When you use the -Server parameter to specify the name of a server to check, the cmdlet tests all mailbox databases on the server simultaneously. If the server is a member of a DAG and has a passive copy of a database, the test is automatically performed against the server that has the active database copy.

Index

A

A records, 108
accepted domains
  defined, 170
  creating, 173, 174
  default, selecting, 174
  identifier, changing, 174
  modifying, 174, 175
  removing, 176
  types, 174
  viewing, 171, 172
actions, for transport rules, 206
Active Directory
  configuring site boundaries, 10
  data store, 14, 15
  database availability groups and, 30
  Edge subscription and, 158
  hub sites, configuring, 110
  site configuration details, viewing, 111
  site-based routing and, 8
  sites, renaming, 111
Active Directory Domain Services (AD DS)
  Edge Transport server role and, 108
  Exchange Server installations and, 2
Active Directory Lightweight Directory Services (AD LDS), 108
Active Manager, 16, 17, 30–34
active monitoring
  managed availability and, 32
  Outlook Web App and, 217
ActiveSync. *See* Exchange ActiveSync
ActiveSyncDevice cmdlets vs. MobileDevice cmdlets, 274
Add cmdlets
  Add-DatabaseAvailabilityGroup-Server, 29, 44, 61
  Add-IPAllowListEntry, 202
  Add-IPBlockListEntry, 204
  Add-IPBlockListProvider, 198, 199, 200, 201
  Add-MailboxDatabaseCopy, 29, 61, 81, 102
Address (A) record, 108
address book policy, 243
address lists, 5
address policies. *See* email address policies
address spaces, Send connectors and, 116, 119, 120
alerting. *See* performance alerting
aliases, missing, 181
anonymous authentication, 237, 238, 239, 240
anti-spam
  described, 189
  agents, 155
  Edge Transport servers and, 154
  IP block lists, 195–201
  Mailbox servers, enabling on, 154
  recipient filtering, 193–195
  sender filtering, 190–193
  updates, configuring, 156
application log, 317
application pools, Outlook Web App and, 220
arbitration mailboxes, 16
architecture changes in Exchange Server 2013, 1, 2, 114
archive mailboxes, 16
ASP.NET Impersonation, 237, 241
attachments
  size limits, 141, 143
  storage location, 18
attempt copy last logs (ACLL), 32
authentication
  See also Forms authentication
  Autodiscover and, 269
  IIS and, 237–241
  for mobile devices, 262
  NTLM, 267

authentication *(continued)*
Outlook Anywhere, 259, 260
POP3 and IMAP4, 253, 254
Receive connectors and, 109, 136, 137
Send connectors and, 114, 119
Windows, 237–239, 253, 254
authoritative domains, 170, 174
autoblocking settings for mobile devices, 297
Autodiscover
described, 268
connection points and, 1
disabling and enabling, 270
health sets, 223
migrating mailboxes and, 13
mobile devices and, 268–271
automated messages, managing mail flow for, 186
automatic distribution, 65
automatic mounting, 97
automatic reseed, 65, 66
availability groups. *See* database availability groups (DAGs)

B

back end transport, 6, 8
back pressure, 169–170
backing up
data, 27
database copies, 17
bandwidth limits, 235
base64 encoding, affect on message size, 141
Basic authentication, 237, 238, 239, 241, 254
bindings
IIS and, 230
POP3 and IMAP4, 252, 253
for Receive connectors, 129
blank senders, 191
blocking messages. *See* message filtering; IP block lists

C

Categorizer, 8, 15
certificates
creating, 233
installing, 234
requesting, 233, 234
trusting, 235
viewing, 232
checkpoint files
message database, 18
queue database, 22
circular logging, 17, 101, 102
Clear-MobileDevice cmdlet
described, 274
syntax and usage, 275
client access
described, 209
layers, 221
throttling, 241–243
Client Access servers
authentication settings for virtual directories, 238
Exchange Server 2013 server roles and, 1
as front end, 2, 114
proxying through, 121
purpose in Exchange Server 2013, 2
Receive connectors and, 128, 129
transport limits, 144
Client front-end connector, 128
Client Proxy connector, 128
client-server architecture in Exchange 2013, 2
client type Receive connectors, 131
cloud service, Exchange Online as, 12
Cluster Continuous Replication (CCR), 28
cluster name objects, 37, 38
clustering, 28, 30, 31
Company filter condition, 180
Compliance Management, journal rules and, 186
conditions, for transport rules, 205
configuration changes, tracking, 343, 344
configuration data type, 14
Connect-Mailbox cmdlet, 77
connection filters. *See* IP block lists
connection limits
domain, 127, 128
simultaneous outbound, 127
website, 235, 236
connection timeout values, 153, 235
ConnectionInactivityTimeout value, 153
connectivity logging
configuring, 314, 315
log fields, list of, 316
properties and fields, 315, 316

connectors. *See* Receive connectors; Send connectors
contacts, Outlook Web App and, 215
Content Filter agent, 155
content indexing, 103–106
continuous replication
changes since previous version, 28
checking status, 58, 59
circular logging, 17
cost of a route
IP site links and, 10
setting for a Send connector, 116
counters, 321–323
CPU utilization, 326, 327
cross-premises routing, 13
current log files
circular logging and, 17
message database, 18, 19
queue database, 22
custom address lists, 5
Custom type Receive connectors, 109, 130
Custom type Send connectors, 109, 117
cutoff balance, throttling client access and, 242

D

data-center coordinator mode, 53
Data Collector Sets, 323
data pages, 19
data storage
Active Directory data store, 14
components, 16
data pages and, 19
data tables and, 19
default location for files, 19
Exchange store, 16
file types, 18
data tables, 19
data types, for replication, 14
database availability groups (DAGs)
described, 28
Active Manager and, 30–34
cmdlets for, 29
creating, 39–43
limit on number of servers, 17
manual network configuration, enabling, 45
membership, managing, 42–45
network cards and, 42
network compression, 51
network encryption, 51
network settings, changing, 49–52
networks, creating, 46, 47
networks, managing, 45–49
permissions, 34, 40
pre-staging and preparing for, 34–39
properties, configuring, 52–55
removing, 54, 55
replication and, 30, 79
restoring operations after member failure, 59–62
routing messages, 8
servers, removing from, 54
witness servers, 34, 35
database copies
activating lagged, 84–86
constraints, 30
creating, 79–81
removing, 94
replication status, monitoring, 90–93
setting replay, truncation, and preference values, 81
status values, list of, 91–93
database management scopes, 65
Database.edb file, 18
databases. *See* mailbox databases
decommissioning Edge Transport servers, 162
Default front-end connector
described, 128
default message size limit, 143
delay notification messages, 127
delayed fan-out, 12, 110
deleted items
recovering, 77, 78
retention periods, 71–75
deleted mailboxes
recovering, 75–77
retention periods, 71–75
delivery groups
boundaries, 8
Exchange versions and, 9
routing destinations and, 6, 110
specifying expansion servers as, 15
Delivery/Relay queue, 20
delivery reports, 305–307
delivery status notification (DSN) codes, 148
Department filter condition, 180

destination delivery groups, 6
device access, 295, 297
dial-up connections, 264
Digest authentication, 237, 238, 239, 241
direct file access, 277–282
Direct Push, 262, 271, 272
discard status, 21
disconnected mailboxes, 75, 76
disk usage, performance alerts for, 327, 328
Dismount-Database cmdlet, 29, 97
dismounting mailbox databases, 95–98
distribution points
  configuring, 245
  properties, 244
DNS lookups, configuring Send connectors, 125, 126
domain data type, 14
Domain Name System (DNS) suffix, Edge Transport server role and, 108
domains
  accepted. *See* accepted domains
  authoritative, 170
  coexistence, 171
  relay, 170
  remote, 186, 187
dynamic distribution groups
  defined, 5
  expansion servers and, 15
  global catalogs and, 14

E

E##00000001.log file, 19
E##.chk file, 18
ECP. *See* Exchange Admin Center (ECP) directory
Edge subscriptions
  Active Directory and, 11
  creating connectors during, 109, 158
  process, 157, 158
  purpose, 156
  removing, 162
  synchronizing, 159
  verifying, 160
  viewing, 159
Edge Transport servers
  anti-spam features, 154
  external relay domains and, 171
  Mailbox servers and, 107
  queues, 20
  role, 2, 108
  subscribing. *See* Edge subscriptions
  viewing Send connectors, 122
EdgeSync service, 156, 159
EHLO start command, 129
E##.log file, 18, 19
email address format options, 178, 180
email address policies
  described, 176
  address book, 243
  creating, 178–180
  custom, variables for, 180
  editing and applying, 182
  priority order, setting, 179, 182
  purpose, 5
  removing, 183
  viewing, 176, 177
.eml message files, 164, 165
encryption
  disabling, 265
  SSL and, 232
enumerating messages, 333
equipment mailboxes, 16
E##Res00001.jrs file, 19
error messages, for probes, 354–356
error tracing, for probes, 356, 357
E##tmp.log file, 18, 19
event logging
  application log, 317–319
  automatic reseed, 66
  Log Parser Studio and, 360
  mailbox database failures, 66
  managed availability and, 33
  monitors and, 352
  probes and, 353, 355
  recovery actions, identifying, 348, 349
  replication and, 30
event sources, list of, 318
Event Viewer, 317–319
EWS. *See* Exchange Web Services (EWS)
Exchange ActiveSync
  described, 261
  configuring URLs and authentication, 247
  disabling and enabling, 263
  IIS and, 263

load balancing, 222
message size limits, 144
resetting, 248
safeguarding sensitive data, 262
synchronization, 262
transport limits, 144–145
web.config file location, 144
Exchange Admin Center (ECP) directory, 223, 248, 249
Exchange Health Manager Service, 33
Exchange Health Manager Worker process, 33
Exchange Management Shell, customizing views, 340
Exchange Online
accepted domains and, 171
capabilities, 12
cmdlets for creating connectors, 140
compliance options, 184
disadvantages of, 12
hybrid deployments, role of Edge Transport servers in, 107
Inbound connectors, 139
vs. on-premises, 12
Outbound connectors, 139
resolving mailbox conflicts with on-premises, 13
Exchange Replication service. *See* replication
Exchange Rich Text Format vs. Transport Neutral Encapsulation Format (TNEF), 189
Exchange Search, 103–106
Exchange Server 2007
Edge Transport servers and, 2
Hub Transport Servers and, 6, 9, 15
journal rules and, 184
online vs. on-premises configurations, 12
Exchange Server 2010
Edge Transport servers and, 2
filter exceptions, 201
global allowed lists, 202
Hub Transport servers and, 6, 9, 15
IP block lists, 197, 200
journal rules and, 184
message filtering, 190, 193
Outlook Web App and, 212
transport dumpster, 21
Exchange Server 2013
Client Access servers and, 1
client-server architecture, 1, 2
continuous replication, 28
data storage, 13
database limits for specific editions, 69
database types, 16
expansion servers, 15
hybrid deployments, 107
installing, 2
Mailbox servers and, 1
mailbox types, 16
managing from touch-capable computers, xiii
on-premises vs. online configuration, 12, 13
organizations and, 2
Outlook versions and, 1
previous editions, changes from, 1, 21, 25, 114, 149
routing and delivery, 9
server roles, 1
server types, 2
upgrading from Standard to Enterprise edition, 69
Exchange store, 16
Exchange Toolbox, 20
Exchange Trusted Subsystem group, 38, 39
Exchange Web Services (EWS)
health sets, 223
limits, 145
message size limit, 144
web.config file location, 145
%ExchangeInstallpath% variable, 144
expansion servers, specifying as a delivery group, 15
expiration time-out
configuring, 127
default, 126
Extended SMTP (ESMTP) sessions, 129
Extensible Storage Engine (ESE), 17, 26, 28
external DNS lookups, configuring, 125
external relay domains, 171, 174
externally secured authentication
Receive connectors and, 137
Send connectors and, 119

F

failover
defined, 29
clusters, 35, 36, 37
vs. switchover, 55
Failover Cluster Manager, 30, 31
fc (file compare) command, 344
file types for Exchange data store, 18
filtering. *See* message filtering; IP block lists
firewalls, 39, 222
Forms authentication, 238, 241
forwarding mailboxes, 16
Front End Transport service
Client Access servers and, 114
incoming messages and, 6
mail transport and, 4
Receive connectors and, 130
routing and, 6
front-end proxying, 121
full-text indexes, 104, 105

G

Get cmdlets
Get-AcceptedDomain, 172
Get-AdSite, 111
Get-AdSiteLink, 112, 113
Get-AutodiscoverVirtualDirectory, 269, 270
Get-ClientAccessServer, 343
Get-DatabaseAvailabilityGroup, 29, 44
Get-DatabaseAvailabilityGroup-Network, 29, 48
Get-EdgeSubscription, 159
Get-EmailAddressPolicy, 177
Get-ExchangeServer, 343
Get-HealthReport, 217, 225, 340
Get-IMAPSettings, 252
Get-IPAllowListEntry, 203
Get-IPBlockListConfig, 199
Get-IPBlockListEntry, 204
Get-IPBlockListProvider, 198
Get-JournalRule, 186
Get-Mailbox, 243
Get-MailboxDatabase, 29, 60, 96, 106
Get-MailboxDatabaseCopyStatus, 29, 62, 93, 102, 106
Get-MailboxServer, 343
Get-MessageTrackingLog, 302, 303
Get-MobileDevice, 274
Get-MobileDeviceMailboxPolicy, 286
Get-MobileDeviceStatistics, 275, 277
Get-MonitoringItemIdentity, 357
Get-OrganizationConfig, 344
Get-OutlookAnywhere, 257, 258
Get-OwaMailboxPolicy, 230
Get-OWAVirtualDirectory, 281
Get-OwaVirtualDirectory, 225
Get-OWAVirtualDirectory, 223, 280, 285
Get-POPSettings, 251
Get-ReceiveConnector, 137
Get-RecipientFilterConfig, 195
Get-RemoteDomain, 186, 187
Get-ResourcePolicy, 343
Get-SendConnector, 124
Get-SenderFilterConfig, 190, 193
Get-ServerHealth, 217, 223, 224, 340, 341
Get-ThrottlingPolicyAssociation, 342
Get-ThrottlingPolicy, 241, 243, 342
Get-TransportRule, 208
Get-TransportService, 344
Get-WorkloadPolicy, 342
global address list, 5
global allowed lists, 202
global block lists, 203
global catalogs, 14
Global Throttling Policy, 342

H

header firewall feature, 109
header size limits, 144
headers, Pickup directory and, 164
health indicators, 343
health sets
ActiveSync, 224
Autodiscover, 223
ECP, 223
Exchange Web Services (EWS), 223

list of, 356, 357
monitors and, 339
OAB, 224
OWA, 218, 219, 224
PublicFolders, 224
RPC, 224
Windows PowerShell, 224
HELO start command, 129
high availability
database options, 28–30
Information Store and, 25
managed availability and, 32
high availability transport (HAT) boundaries, shadow redundancy and, 149
hops, specifying maximum for Receive connectors, 133
HTTP
default port settings, 251
optimizing protocol logging for, 311–313
protocol log files, 313, 314
redirection role service, 237
server, controlling access to, 237–241
hub sites, 110, 111
Hub Transport servers, 2, 6, 9
hybrid deployments
Mailbox servers and Edge Transport servers, 107
message routing and delivery, 107
on-premises and Exchange Online, connectors for, 139
hybrid organizations
accepted domains and, 171
journal rules and, 184

I

idle user sessions, disconnecting, 235
IIS
authentication, 237–241
bindings, 230–232
Exchange ActiveSync and, 263
Log Parser Studio and, 360
restarting, 225
troubleshooting, 225, 226
web applications and, 221, 222
IMAP4 (Internet Message Access Protocol 4)
authentication, 253, 254
bindings, 252, 253
configuring, 249–250
default port settings, 251
enabling, 250–251
message retrieval settings, 256, 257
Inbound connectors, 139
incoming messages, Mailbox Transport Delivery service and, 6
incremental resync, 87
indexing, 103–106
information store, 16, 17, 25–34, 63
installing Exchange Server 2013, 2
Integrated Windows Authentication, Receive connectors and, 137
internal DNS lookups, configuring, 125
Internal Relay Domain accepted domain type, 174
internal relay domains, 171
Internal type Receive connectors, 109, 130
Internal type Send connectors, 109, 117
Internet Message Access Protocol 4 (IMAP4). *See* IMAP4 (Internet Message Access Protocol 4)
Internet type Receive connectors, 131
Internet type Send connectors, 117
interoperability with previous versions, site-based routing and, 8
Invoke-MonitoringProbe cmdlet, 358
IP Allow List, enabling, 202
IP block list services
enabling and disabling, 200
priority order, 198, 200
return status codes for, 196, 197, 198
IP block lists
applying, 196, 197, 198
enabling, 203
error messages, 200, 201
exceptions, 201, 202
IP site links, 10, 112, 113
IPv4 addresses, specifying for Receive connectors, 131
IPv6 addresses, specifying for Receive connectors, 131

J

journal rules
configuring, 184
creating, 185, 186
managing, 186

L

lagged database copies, activating, 84–86
least-cost routing path
default behavior, changing, 113
determining, 109
message relay, 110
legacy Edge Transport servers. *See* Edge Transport servers
Lightweight Directory Access Protocol (LDAP) queries, 14
line length, setting maximum, 189
linked mailboxes, 5, 16
load balancing
Exchange ActiveSync, 222
excluding a mailbox from provisioning, 66
mailbox databases, 64, 65
Outlook Web App, 222
local area networks (LANs), as sites, 10
Local Continuous Replication (LCR), 28
log files, 17–20
See also transaction logs
current active, 22, 23
location, 23, 26
message tracking, 300–305
reserve, 22, 23
secondary, 22
Log Parser Studio, 359–361
logging
See also event logging
circular, 17, 101, 102
connectivity, 314–316
protocol, 307–314

M

mailbox databases
described, 16
active vs. passive, 30, 63, 78
associated default files, 26
backing up data, 27
copies, creating, 79–81
copies, removing passive, 94
copy status values, 91–93
creating, 66–67
default, 64
deleting, 103
high availability options, 28
limits on specific editions of Exchange, 69
load balancing, 64, 65
maintenance intervals, setting, 98–100
mounting and dismounting, 95–98
moving, 100–102
naming, 26
primary vs. secondary, 78
recovering databases, 95, 96
recovering items deleted from, 77, 78
recovering mailboxes deleted from, 75–77
renaming, 102, 103
replication status, monitoring, 90–93
replication, suspending and resuming, 83, 84
retention periods, setting, 71–75
status, 96
storage limits, 26, 72, 73
system requirements, 17
updating copies, 87–90
viewing file and folder paths, 100
mailbox delivery queues, 329
mailbox-enabled recipients, 5
mailbox policies for mobile devices
assigning, 293
creating, 288–292
optimizing, 291
options, 288–290
purpose, 263
removing, 294
viewing, 285–287
mailbox policies for Outlook Web App, 216, 217, 229, 230
mailbox provisioning load balancer. *See* load balancing
Mailbox servers
anti-spam features, enabling, 154
authentication settings for virtual directories, 239
as back end, 2, 114
connectors and, 109

database limit and maximum size, 17
Edge Transport servers and, 107, 114
Exchange Server 2013 server roles and, 1
external DNS lookups, configuring, 125
internal DNS lookups, configuring, 126
queues, 20
Receive connectors and, 128, 129
routing and, 8, 9
Mailbox Transport Delivery service, incoming messages and, 6
Mailbox Transport service, 6, 8, 9
Mailbox Transport Submission service, outgoing messages and, 6
mailbox types, 16
mail-enabled recipients, 5
mail flow, testing, 345, 346
mail protocols, 108
Mail.que file, 22
maintenance intervals, setting, 98–100
managed availability, 32–34, 337, 338
Managed Store. *See* information store
MaxAllowedContentLength value, 145, 146
MaxDocumentDataSize key, ActiveSync limits and, 145
MaxReceivedMessageSize value, 145, 146
MaxRequestLength value, 145, 146
memory usage, performance alerting and, 323–326
MessageExpirationTimeout value, 153
message filtering
described, 189
blocking domains and senders, 192
Edge Transport servers and, 154
IP block lists. *See* IP block lists
preventing on internal servers, 205
by recipient, 193–195
by sender, 190–193
message tracking
configuring logs, 300, 301
delivery reports, 305–307
log fields, list of, 304
purpose, 299
reviewing logs manually, 303–305
searching logs, 302, 303
settings, 300
messages
header limits, 141
Pickup directory limits, 167
processing speed, 166, 167
queueing, 11, 20, 21, 22
receive size limits, 141, 142
routing, 4, 6, 8–13
send size limits, 141, 143
throttling, 168, 169
messaging networks, 45, 50
Microsoft Exchange Online. *See* Exchange Online
Microsoft Exchange Server 2013. *See* Exchange Server 2013
Microsoft.Exchange.Store.Service.exe. *See* Information Store
Microsoft Exchange Transport service (MSExchangeTransport.exe), 23
Microsoft Exchange Writer, 17
Microsoft Outlook. *See* Outlook 2013
Microsoft-Server-ActiveSync directory, 224
migrating mailboxes from online to on-premises, 13
MIME types, 277, 278, 280, 281
mobile devices
Autodiscover and, 268–271
direct file access, 277–282
Direct Push and, 271, 272
Exchange ActiveSync and, 262, 263
mailbox policies, 285–295
managing access, 295–297
Outlook Anywhere and, 263–267
password recovery, 276–277
remote device wipe and, 272–275
remote file access, 282
WebReady Document Viewing, 283–284
monitor engine, 338
monitors, 352
Mount-Database cmdlet, 97
mounting mailbox databases, 57, 58, 68, 95–98
Move cmdlets
Move-ActiveMailboxDatabase, 29, 57
Move-DatabasePath, 29, 100
MSExchangeDagMgmt.exe. *See* database availability groups (DAGs)
MSExchangeHMHost.exe, 33

MSExchangeHMWorker.exe, 33
MSExchangeRepl.exe. *See* replication service
msExchMDBAvailabilityGroup object, 35, 36
multihomed servers, 230
multimaster replication, 14
Multipurpose Internet Mail Extensions (MIME) values, 277, 278
multivalued properties, shorthand for, 202
MX records, 108, 118

N

naming Exchange organizations, 2
NDR journaling mailbox, 184, 185
networks, database availability group, 45–52
New cmdlets
    New-AcceptedDomain, 174
    New-AutodiscoverVirtualDirectory, 270
    New-DatabaseAvailabilityGroup, 29, 41
    New-DatabaseAvailabilityGroupNetwork, 29, 49
    New-EcpVirtualDirectory, 225
    New-EdgeSubscription, 158
    New-EmailAddressPolicy cmdlet, 181
    New-JournalRule, 186
    New-MailboxDatabase, 29, 70
    New-MobileDeviceMailboxPolicy, 290, 291
    New-OwaMailboxPolicy, 230
    New-OwaVirtualDirectory, 225
    New-OWAVirtualDirectory, 223
    New-ReceiveConnector cmdlet, 133, 134, 143, 144
    New-RemoteDomain cmdlet, 187
    New-ResourcePolicy, 343
    New-SendConnector cmdlet, 121, 143
    New-ThrottlingPolicy, 241, 243, 342
    New-TransportRule, 208
    New-WorkloadManagementPolicy, 343
    New-WorkloadPolicy, 342
New Mailbox Database Wizard, 66
non-delivery of messages, Mail Transport service and, 8
non-delivery reports, 127, 141, 144, 148
nonpersistent queues, 20
non-SMTP connectors, Replay directory and, 164
Ntds.dit file, 14

O

OAB directory
    health sets, 224
    resetting, 245
    web distribution points, 244
object-based storage, 19
offline address book (OAB), 64, 70, 71
Offline mode, Outlook Web App and, 215
On-Premises connector type, 139
organizations
    Exchange environments and, 2
    hybrid deployments, 107
    recipients, managing, 5
Outbound Connection Failure Retry Interval (Seconds), 127
outbound connections, 127, 128
Outbound connectors, 139
Outbound Proxy front-end connector
    default message size limit, 143
    purpose, 128
outgoing messages, Mailbox Transport Submission service and, 6
Outlook 2013
    Outlook Anywhere and, 264
    recovering deleted items, 77
Outlook Anywhere
    described, 264
    configuring, 257–260, 265–269
    connections and, 3
    RPC over HTTP and, 264, 266
    vs. remote mail, 264
Outlook Web Access
    disabling and enabling, 263
    mailboxes, accessing, 213
    public folders, accessing, 213
Outlook Web App
    described, 210, 261
    Active Monitoring and, 217
    application pools, configuring, 220
    configuration requirements, 210
    contacts, 215
    disabling, 216

features, 211, 212, 214
health sets, 217, 218, 219
limits, 146
load balancing, 222
mailboxes, accessing, 212
mailbox policies, 216, 217, 229, 230
message size limit, 144
Offline mode, 215
password recovery, 276
remote device wipe and, 273
segmentation features, 226–230
troubleshooting, 217–221
versions, 210, 211
web.config file location, 146
overlapping recycling of worker processes, 147
OWA. *See* Outlook Web App
OWA directory
configuring URLs and authentication, 246
health sets, 224
resetting, 247

P

page headers, 19
Partner connection type, 140
Partner type Receive connectors, 131
Partner type Send connectors
described, 117
header firewall feature and, 109
password recovery for mobile devices, 276
performance alerting, 323–328
performance monitoring, 321–323
performance objects, 321–323
persistent queues, 20
Pickup directory
header fields and, 164
message routing and delivery, 163
messaging limits, 167
moving, 165
polling interval, 164
Transport servers and, 164
poison message queue, 20, 329
policies
address book, 243
email address, 176–183
Global Throttling, 342
mailbox, for mobile devices, 285–294
mailbox, for Outlook Web Access, 216, 229, 230
mailbox, for Outlook Web App, 217
remote domains, 186
retention, 5, 243
polling interval
maximum messages processed, 166
Pickup directory, 164
Replay directory, 165
POP3 (Post Office Protocol 3)
authentication, 253, 254
bindings, 252, 253
configuring, 249
connection settings, 254–256
default port settings, 251
enabling, 250, 251
message retrieval settings, 256, 257
port settings for messaging protocols, 251
postmaster address, 147, 148
Post Office Protocol 3 (POP3). *See* POP3 (Post Office Protocol 3)
PowerShell. *See* Windows PowerShell
preference values, 80, 81–83, 91
preprovisioned log files
message database, 18
queue database, 22
pre-staging cluster name objects, 37
primary data files, 18, 19, 22
primary role holders, 31, 32
probe engine, 338, 339, 354
probes
definitions, 353
error messages, 354–356
monitors and, 352
tracing errors, 356–358
properties
for Receive connectors, 136
for Send connectors, 123
Protocol Analysis agent, 155
protocol logging
configuring, 307–309
optimizing for HTTP, 311–313
properties and fields, 310, 311
purpose, 307
receive connector parameters, list of, 309
send connector parameters, list of, 309
working with log files, 313, 314

provisioning with Autodiscover, 268
proxying through Client Access servers, 121
proxy settings, 266
Public directory, 224
public folder mailboxes, 16
public folders, Send As permission and, 5

Q

queueing messages. *See* messages
queues
  deleting, 335
  filtering messages, 333, 334
  forcing connections to, 334
  states, 331
  summaries, 331
  suspending and resuming, 334
  types, 328–330
  viewing, 330, 331, 333
Queue Viewer, 20, 330, 331, 332
quorum, 35
quorum resource, 31, 37
quota notifications, 73, 74

R

Receive connectors
  described, 109
  bindings, 129
  creating, 128, 129, 130, 131, 133
  defaults created to enable mail flow, 128
  disabling and enabling, 136
  header size limits, 144
  managing, 135
  maximum hops, specifying, 133
  permissions, 137
  properties, changing, 136
  protocol logging, parameters for, 309
  recipient limits, 144
  removing, 136
  security, 136
  specifying IP addresses, 131
  transport limits, 143
  types of, 130
  viewing, 135
recharge rate, client access and, 242
Recipient Filter agent, 155
recipient limits, 144
recipients
  defined, 5
  filtering, 179, 180, 193, 194
  limits, 141
  managing, 5
  resolution of, 8
recovering mailbox databases, 95, 96
recovery actions, 348–350
recovery response, levels, 339
redirecting users to alternate URLs, 236, 237
relay domains, 170
relaying messages, 110
relay queues, 329
remote delivery queues, 20, 329
remote device wipe, 272–275
remote domains, 186–189
  creating, 187
  message options for, 187, 188, 189
  policies, 186
  removing, 189
  settings, 186
  viewing, 186
remote file access, 282
remote mail vs. Outlook Anywhere, 264
remote network settings, specifying, 131
Remove cmdlets
  Remove-AcceptedDomain, 176
  Remove-AutodiscoverVirtual-Directory, 270, 271
  Remove-DatabaseAvailabilityGroup, 29, 55
  Remove-DatabaseAvailability-GroupNetwork, 29, 49
  Remove-DatabaseAvailability-GroupServer, 29, 45, 61
  Remove-EdgeSubscription, 162
  Remove-EmailAddressPolicy, 183
  Remove-IPAllowListEntry, 203
  Remove-IPBlockList, 204
  Remove-IPBlockListProvider, 198
  Remove-JournalRule, 186
  Remove-MailboxDatabase, 29, 103
  Remove-MailboxDatabaseCopy, 29, 60, 61, 94
  Remove-OwaVirtualDirectory, 225
  Remove-OWAVirtualDirectory, 223

Remove-ReceiveConnector, 139
Remove-RemoteDomain, 189
Remove-ResourcePolicy, 343
Remove-SendConnector, 125
Remove-ThrottlingPolicy, 241, 243, 342
Remove-TransportRule, 208
Remove-WorkloadPolicy, 342
renaming Active Directory sites, 111
Replay directory, 163–166
message routing and delivery, 163
moving, 166
non-SMTP connectors and, 164
polling interval, 165
replay lag time, 81–83, 91
replaying logs into lagged copies, 82, 83
replication
automatic reseed and, 65
database availability groups and, 79
disabling, 46, 47
event logging, 30
monitoring status, 90–93
multimaster, 14
networks, 45, 50
seeding and, 87
service, 16, 17
suspending, 81, 83, 84
testing, 346, 347
reseeding databases, 65, 66, 87
reserve log files
message database, 19
queue database, 22
resource thresholds, 343
responder engine, 338, 339
responders, identifying, 350
restarting websites, 244
Restore-DatabaseAvailabilityGroup cmdlet, 29, 54
restoring databases, 17, 65, 66
result types, 358
Resume-MailboxDatabaseCopy cmdlet, 29, 84
retention periods, 71–75
retention policies, 5, 243
retrying unsuccessful outbound connections, 127
return status codes for IP block list services, 196–198
role assignment policy, 243
roles, primary vs. secondary Active Managers, 31, 32
room mailboxes, 16
routing
cost of a route, 10
cross-premises, 13
delivery groups and, 6, 9
distribution group expansion servers, 15
Front End Transport service and, 4, 6
incoming messages, 6
multi-recipient messages, 12
site-based vs. group based, 8
routing destinations, delivery groups and, 110
routing headers, 109
row offsets, 19
RPC directory, 224, 258, 259
RPC MAPI, Transport service and, 6
RPC over HTTP. *See* Outlook Anywhere
RPC over TCP/IP, Outlook Anywhere and, 267
RPC sessions, Mailbox servers and, 3

S

SafeList aggregation feature, 193
Safety Net
hold time, 154
parameters, list of, 154
queues, 20, 21, 22, 150, 153, 330
SAN (storage area network) implementation, 27
schema data type, 14
searching
content indexing and, 103–105
tracking logs, 302, 303
secondary log files
mailbox database, 19
queue database, 22
secondary role holders, 32
security
Exchange ActiveSync safeguards, 262
groups, specifying for Receive connectors, 137
header firewall feature, 109
mobile device mailbox policies, 263

security (*continued*)
Outlook Anywhere and, 264, 267
Pickup directory permissions, 165
Receive connectors and, 136
remote device wipe, 272–275
Replay directory permissions, 165
seeding databases, 87–90
segmentation, 226
segmentation features
enabling and disabling, 228, 229
listed, 227
Send As permission, 5
Send connectors
described, 108
address spaces and, 119, 120
creating, 114, 115, 116, 117
disabling, 123
DNS lookups, 125, 126
enabling, 123
for internal mail flow, 114
for Internet mail flow, 114, 115, 116
limits, 126, 127
linking to specific Receive connectors, 115
managing, 122, 123
maximum size for messages, 116, 121
properties, changing, 123
protocol logging, parameters for, 309
removing, 123
setting cost of, 116
transport limits, 143
types, 117
viewing, 122, 123
Sender Filter agent, 155
sender filtering, 190, 192, 193
Sender ID agent, 155
server health
tracking, 337–342
transition states, list of, 33, 34
service health, testing, 344
Set cmdlets
Set-AcceptedDomain, 175
Set-ActiveSyncVirtualDirectory, 239, 282
Set-AdSite, 111, 112
Set-AdSiteLink, 113
Set-AutodiscoverVirtualDirectory, 239, 269, 271
Set-CASMailbox, 230, 294
Set-ContentFilterConfig, 201, 205
Set-DatabaseAvailabilityGroup, 29, 35, 37, 51, 52, 53, 82
Set-DatabaseAvailabilityGroupNetwork, 29, 51
Set-EcpVirtualDirectory, 225, 239
Set-EmailAddressPolicy, 183
Set-ExchangeServer, 343
Set-IMAPSettings, 252
Set-IPBlockListConfig, 199
Set-IPBlockListProvider, 198, 200, 201, 205
Set-IPBlockListProvidersConfig, 199, 202
Set-JournalRule, 186
Set-MailboxDatabase, 29, 65, 71, 74, 75, 98, 99, 100, 102, 103
Set-MailboxDatabaseCopy, 29, 82, 154
Set-MailboxServer, 58
Set-MailboxTransportService, 315
Set-MobileDeviceMailboxPolicy, 291, 292, 293
Set-OabVirtualDirectory, 239
Set-OutlookAnywhere, 260
Set-OwaMailboxPolicy, 230
Set-OWAVirtualDirectory, 225, 239, 280, 281, 284, 285
Set-POPSettings, 251
Set-PowerShellVirtualDirectory, 239
Set-ReceiveConnector, 138, 143, 144, 153, 169
Set-RecipientFilterConfig, 195
Set-RemoteDomai, 188
Set-ResourcePolicy, 343
Set-SendConnector, 124, 143, 153, 168
Set-SenderFilterConfig, 190, 191, 192
Set-SenderIdConfig, 205
Set-Service, 251
Set-ThrottlingPolicy, 241, 342
Set-ThrottlingPolicyAssociation, 342
Set-TransportConfig, 143, 151, 152, 154
Set-TransportRule, 208
Set-TransportServer, 147

Set-TransportService, 153, 154, 155, 164, 166, 167, 168, 300, 315
Set-TransportService cmdlet, 165
Set-WebServicesVirtualDirectory, 239
Set-WorkloadPolicy, 342
shadow redundancy
process, 149, 150
purpose, 148
queue, 20, 21, 329
shared mailboxes, 16
sharing policy, 243
Simple Mail Transfer Protocol (SMTP). *See* SMTP
site-based routing, 8
site bindings, 230–232
site links
described, 10
configuration information, viewing, 112
Exchange-specific costs, setting, 113
maximum size for messaages, 113
relaying messages across, 11
site membership, for Exchange servers, 9
slow networks, defining, 267
smart hosts
defined, 114
adding to Send connectors, 118
authentication and, 119
SMTP
address spaces, 119, 120
connectors, 108
default port settings, 251
delivering messages, 6
primary address, 8
routing messages, 4
Send connectors, for Internet mail flow, 115, 116
standard vs. extended sessions, 129
soft-deleted mailboxes, 75
source servers, 120
source transport servers, 6
spam confidence levels, 191
spam filtering, 190, 191, 192, 193
SSL
enabling on websites, 232–235
Outlook Anywhere and, 266
Standby Continuous Replication (SCR), 28
Start-DatabaseAvailabilityGroup cmdlet, 29, 53
starting websites, 244
State Or Province filter condition, 179
State values, list of, 341
Stop-DatabaseAvailabilityGroup cmdlet, 29, 54
stopping websites, 244
storage area network (SAN) implementation, 27
storage limits for mailbox databases, 72, 73
submission queue, 20, 328
subscribing Edge Transport servers, 156, 158
Suspend-MailboxDatabaseCopy cmdlet, 29, 84
switching over servers and databases, 55–58
switchover
defined, 29
vs. failover, 55
synchronizing
Edge subscriptions, 158
mobile devices, 271
preventing, 295
system services, 320, 321

T

temporary data files
message database, 18
queue databasee, 22
temporary log (E##tmp.log) file, 19
Test cmdlets
Test-ActiveSyncConnectivity, 347
Test-ArchiveConnectivity, 347
Test-AssistantHealth, 347
Test-CalendarConnectivity, 347
Test-EcpConnectivity, 235, 347
Test-EdgeSynchronization, 160, 161, 347
Test-ExchangeSearch, 106, 347
Test-FederationTrust, 347
Test-FederationTrustCertificate, 347
Test-IMAPConnectivity, 252, 347
Test-IPAllowListProvider, 347

Test cmdlets *(continued)*
Test-IPBlockListProvider, 199, 347
Test-IRMConfiguration, 347
Test-Mailflow, 345
Test-MAPIConnectivity, 252, 347
Test-MRSHealth, 347
Test-OAuthConnectivity, 347
Test-OutlookConnectivity, 347
Test-OutlookWebServices, 235, 348
Test-OWAConnectivity, 217, 223, 235, 348
Test-POPConnectivity, 252, 348
Test-PowerShellConnectivity, 348
Test-ReplicationHealth, 58, 59, 102, 346, 347
Test-SenderId, 348
Test-ServiceHealth, 344
Test-SmtpConnectivity, 348
Test-UMConnectivity, 348
Test-WebServicesConnectivity, 348
throttling
client access, 241–243
messages, 168, 169
users, 342, 343
workloads, 342, 343
time-out values, setting, 235, 236
TLS connections, 253, 254
Tmp.edb file, 18, 22
TNEF (Transport Neutral Encapsulation Format), remote domains and, 189
touch-capable computers, managing Exchange Server 2013 using, xiii
tracking logs
event fields, 304
importing, 305
purpose, 300
searching, 302, 303
tracking messages. *See* message tracking
transaction logs, 17, 18, 19, 22
Transient Failure Retry Attempts, 127
Transient Failure Retry Interval (Minutes), 127
transition state, 33, 34
transport dumpster queue, 330
Transport Layer Security, 136
transport limits
described, 141
connector, 143, 144
organizational, 142, 143
retrying unsuccessful connections, 127
server, 144, 145, 146
Transport Neutral Encapsulation Format (TNEF) message data, remote domains and, 189
transport rules
actions, 206, 207
conditions, 205, 207
creating, 206, 207
exceptions, 206, 208
managing, 208
transport servers
destination delivery group and, 6
DNS lookups, 125
Pickup directory and, 164
queues, 20
Send connectors and, 114
Transport service
back-end transport and, 6–8
Mailbox servers and, 114
Receive connectors and, 130
transporting messages. *See* Front End Transport service; Mailbox Transport services; routing
TRN00000001.log file, 22
Trn.chk file, 22
Trn.log file, 22
TRNRes00001.jrs file, 22
Trntmp.log file, 22
troubleshooting steps, 348
truncation lag time, 81–83

U

Unified Messaging servers, 2
unreachable queue, 20, 330
Update-EmailAddressPolicy cmdlet, 181, 182
Update-MailboxDatabaseCopy cmdlet, 29, 89, 90, 106
upgrading to Exchange Server 2013 Enterprise edition, 69, 70
user mailboxes, 16
user throttling, 342, 343

V

virtual directories
authentication on Client Access servers, 238
authentication on Mailbox servers, 239
recreating, 225
web applications and, 222–226
virtual directories for Autodiscover, resetting, 269

W

web applications
IIS and, 221, 222
root virtual directory, 222
web distribution points
configuring, 245
properties, 244
web.config file location
Exchange ActiveSync, 144
Exchange Web Services, 145
Outlook Web App, 146
WebReady Document Viewing, 283–285
website limits, 236
websites, starting and stopping, 244
wide area network (WAN) links, as site boundaries, 10
Windows authentication, 237, 238, 239, 253, 254
Windows Failover Clustering, 28, 30, 36
Windows Firewall, witness server and, 39
Windows PowerShell
directory, 224
throttling, 242
witness servers, 34, 35, 36, 40, 52, 53
working files, 18, 19
workload classifications, 342
workload throttling, 342, 343
wrapping text, 189

X

X.509 certificates, 232
X-headers, 109