

The Administrator's Reference

2 Books
in 1

Exchange Server 2016

Super-sized desktop reference

Combines 2 training guides in 1 volume

Expert advice at your fingertips

@techjob

William R. Stanek
Author

William R. Stanek Jr.
Contributor

Acknowledgments

To my readers—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 everyone I've worked with at Microsoft—thanks for the many years of support and for helping out in ways both large and small.

Special thanks to my son Will for his extensive contributions to this book. You've made many contributions previously, but now I can finally give you the cover credit you've earned and deserved for so long.

—William R. Stanek

Exchange Server 2016: The Administrator's Reference

William R. Stanek
Author & Series Editor

William R. Stanek, Jr.
Contributor

Exchange Server 2016

The Administrator's Reference

Published by Stanek & Associates
PO Box 362, East Olympia, WA, 98540-0362
www.williamrstanek.com

© 2017 William R. Stanek. Seattle, Washington.
All rights reserved.

No part of this book may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted by Sections 107 or 108 of the 1976 United States Copyright Act, without the prior written permission of the publisher. Requests to the publisher for permission should be sent to the address listed previously.

Stanek & Associates is a trademark of Stanek & Associates and/or its affiliates. All other marks are the property of their respective owners. No association with any real company, organization, person or other named element is intended or should be inferred through use of company names, web site addresses or screens.

This book expresses the views and opinions of the author. The information contained in this book is provided without any express, statutory or implied warranties.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER AND THE AUTHOR MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES OR PROMOTIONAL MATERIALS. THE ADVICE AND DISCUSSION IN THIS BOOK MAY NOT BE SUITABLE FOR EVERY SITUATION. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING PROFESSIONAL SERVICES AND THAT SHOULD PROFESSIONAL ASSISTANCE BE REQUIRED THE SERVICES OF A COMPETENT PROFESSIONAL SHOULD BE SOUGHT. NEITHER THE PUBLISHERS, AUTHORS, RESELLERS NOR DISTRIBUTORS SHALL BE HELD LIABLE FOR ANY DAMAGES CAUSED OR ALLEGED TO BE CAUSED EITHER DIRECTLY OR INDIRECTLY HEREFROM. THE REFERENCE OF AN ORGANIZATION OR WEBSITE AS A SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE PUBLISHER OR THE AUTHOR ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR THE RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT WEBSITES LISTED IN THIS BOOK MAY NOT BE

AVAILABLE OR MAY HAVE CHANGED SINCE THIS WORK WAS WRITTEN.

Stanek & Associates publishes in a variety of formats, including print, electronic and by print-on-demand. Some materials included with standard print editions may not be included in electronic or print-on-demand editions or vice versa.

Country of First Publication: United States of America.

Cover Design: Creative Designs Ltd.

Editorial Development: Andover Publishing Solutions

Technical Review: L & L Technical Content Services

You can provide feedback related to this book by emailing the author at williamstanek@aol.com. Please use the name of the book as the subject line.

Version: 2.1.1.3b

Note I may periodically update this text and the version number shown above will let you know which version you are working with. If there's a specific feature you'd like me to write about in an update, message me on Facebook (<http://facebook.com/williamstanekauthor>). Please keep in mind readership of this book determines how much time I can dedicate to it.

Table of Contents

[How to Use This Guide](#)

[Print Readers](#)

[Digital Book Readers](#)

[Support Information](#)

[Conventions & Features](#)

[Share & Stay in Touch](#)

[Chapter 1. Welcome to Exchange 2016](#)

[Getting Started with Exchange Admin Center](#)

[Navigating Exchange Admin Center Options](#)

[Accessing Exchange Admin Center](#)

[Authenticating and Proxying Connections](#)

[Getting Started with Exchange Management Shell](#)

[Running and Using Cmdlets](#)

[Running and Using Other Commands and Utilities](#)

[Using Cmdlet Parameters and Errors](#)

[Using Cmdlet Aliases](#)

[Working with Exchange Management Shell](#)

[Starting Exchange Management Shell](#)

[Using Exchange Cmdlets](#)

[Working with Object Sets and Redirecting Output](#)

[Chapter 2. Working with Exchange Online](#)

[Getting Started with Exchange Online](#)

[Navigating Exchange Online Services](#)

[Understanding Office 365 Licensing](#)

[Using Windows PowerShell with Exchange Online](#)

[Getting Started with Windows PowerShell](#)

[Understanding the Default Working Environment](#)

[Learning About Cmdlets and Functions](#)

[Connecting to Exchange Online Using PowerShell](#)

[Exploring How the Shell Uses Remote Sessions](#)

[Establishing Remote Sessions](#)

[Using an Interactive Remote Session](#)

[Creating and Importing a Remote Session](#)

[Connecting to Windows Azure](#)

[Cmdlets for Windows Azure Active Directory](#)

[Working with Exchange Online Cmdlets](#)

[Cmdlets Specific to Exchange Online](#)

[Working with Exchange Online Cmdlets](#)

Chapter 3. Getting Started with Users and Contacts

Working with Users and Contacts

How Email Routing Works: The Essentials

Managing Recipients: The Fundamentals

Finding Existing Mailboxes, Contacts, And Groups

Finding Synced, Unlicensed, Inactive, and Blocked Users

Chapter 4. Managing Users

Creating Mailbox-Enabled and Mail-Enabled User Accounts

Working with Logon Names and Passwords

Mail-Enabling New User Accounts

Mail-Enabling Existing User Accounts

Managing Mail-Enabled User Accounts

Creating Domain User Accounts with Mailboxes

Creating Online User Accounts with Mailboxes

Adding Mailboxes to Existing Domain User Accounts

Setting or Changing the Common Name and Logon Name for Domain User Accounts

Setting or Changing Contact Information for User Accounts

Changing Logon ID or Logon Domain for Online Users

Changing a User's Exchange Server Alias and Display Name

Adding, Changing, and Removing Email and Other Addresses

Setting a Default Reply Address for a User Account

Changing A User's Web, Wireless Service, And Protocol Options

Requiring Domain User Accounts to Change Passwords

Deleting Mailboxes from User Accounts

Deleting User Accounts and Their Mailboxes

Chapter 5. Managing Contacts

Creating Mail-Enabled Contacts

Setting or Changing a Contact's Name and Alias

Setting Additional Directory Information for Contacts

Changing Email Addresses Associated with Contacts

Disabling Contacts and Removing Exchange Attributes

Deleting Contacts

Chapter 6. Adding Special-Purpose Mailboxes

Using Room and Equipment Mailboxes

Adding Room Mailboxes

Adding Equipment Mailboxes

Adding Linked Mailboxes

Working with Archive Mailboxes

Adding In-Place Archives

Adding Online Archives

Managing Archive Settings

[Adding Arbitration Mailboxes](#)

[Adding Discovery Mailboxes](#)

[Adding Shared Mailboxes](#)

[Adding Public Folder Mailboxes](#)

[Chapter 7. Managing Mailboxes](#)

[Managing Mailboxes: The Essentials](#)

[Viewing Current Mailbox Size, Message Count, and Last Logon](#)

[Configuring Apps for Mailboxes](#)

[Hiding Mailboxes from Address Lists](#)

[Defining Custom Mailbox Attributes for Address Lists](#)

[Restoring On-Premises Users and Mailboxes](#)

[Restoring Online Users and Mailboxes](#)

[Repairing Mailboxes](#)

[Moving Mailboxes](#)

[Importing and Exporting Mail Data](#)

[Performing On-Premises Mailboxes Moves and Migrations](#)

[Performing On-Premises Mailbox Moves](#)

[Moving Mailboxes Within a Single Forest](#)

[Moving Mailboxes Between Forests](#)

[Managing Delivery Restrictions, Permissions, and Storage Limits](#)

[Setting Message Size Restrictions for Contacts](#)

[Setting Message Size Restrictions on Delivery to and from Individual Mailboxes](#)

[Setting Send and Receive Restrictions for Contacts](#)

[Setting Message Send and Receive Restrictions on Individual Mailboxes](#)

[Permitting Others to Access a Mailbox](#)

[Forwarding Email to a New Address](#)

[Setting Storage Restrictions on Mailbox and Archives](#)

[Setting Deleted Item Retention Time on Individual Mailboxes](#)

[Chapter 8. Managing Groups](#)

[Using Security and Distribution Groups](#)

[Group Types, Scope, And Identifiers](#)

[When to Use Security and Standard Distribution Groups](#)

[When to Use Dynamic Distribution Groups](#)

[Working with Security and Standard Distribution Groups](#)

[Group Naming Policy](#)

[Understanding Group Naming Policy](#)

[Defining Group Naming Policy for Your Organization](#)

[Defining Blocked Words in Group Naming Policy](#)

[Creating Security and Standard Distribution Groups](#)

[Creating a New Group](#)

[Mail-Enabling Universal Security Groups](#)

[Assigning and Removing Membership for Individual Users, Groups, and Contacts](#)

[Adding and Removing Managers](#)

[Configuring Member Restrictions and Moderation](#)

[Working with Dynamic Distribution Groups](#)

[Creating Dynamic Distribution Groups](#)

[Changing Query Filters and Filter Conditions](#)

[Designating an Expansion Server](#)

[Modifying Dynamic Distribution Groups Using Cmdlets](#)

[Previewing Dynamic Distribution Group Membership](#)

[Other Essential Tasks for Managing Groups](#)

[Changing a Group's Name Information](#)

[Changing, Adding, or Deleting a Group's Email Addresses](#)

[Hiding Groups from Exchange Address Lists](#)

[Setting Usage Restrictions on Groups](#)

[Creating Moderated Groups](#)

[Deleting Groups](#)

[Chapter 9. Managing Addresses Online and Offline](#)

[Managing Online Address Lists](#)

[Using Default Address Lists](#)

[Using Address Book Policies](#)

[Creating and Applying New Address Lists](#)

[Updating Address List Configuration and Membership Throughout the Domain](#)

[Previewing and Editing Address Lists](#)

[Configuring Clients to Use Address Lists](#)

[Renaming and Deleting Address Lists](#)

[Managing Offline Address Books](#)

[Creating Offline Address Books](#)

[Configuring Clients to Use an Offline Address Book](#)

[Setting the Default Offline Address Book](#)

[Changing Offline Address Book Properties](#)

[Designating OAB Generation Servers and Schedules](#)

[Rebuilding the OAB](#)

[Deleting Offline Address Books](#)

[Chapter 10. Configuring Exchange Clients](#)

[Mastering Outlook Web App essentials](#)

[Getting started with Outlook Web App](#)

[Connecting to Mailboxes and Public Folder Data Over the Web](#)

[Working with Outlook Web App](#)

[Enabling and Disabling Web Access for Users](#)

[Configuring Mail Support for Outlook](#)

[Understanding Address Lists, Offline Address Books, and Autodiscover](#)

[Configuring Outlook for the First Time](#)

[First-Time Configuration: Connecting to Exchange Server](#)

[First-Time Configuration: Connecting to Internet Email Servers](#)

[Configuring Outlook for Exchange](#)

[Adding Internet Mail Accounts to Outlook](#)

[Repairing and Changing Outlook Mail Accounts](#)

[Leaving Mail on the Server with POP3](#)

[Checking Private and Public Folders with IMAP4 and UNIX Mail Servers](#)

[Managing the Exchange Configuration in Outlook](#)

[Managing Delivery and Processing Email Messages](#)

[Using Server Mailboxes](#)

[Using Personal Folders](#)

[Repairing .pst data files](#)

[Repairing .ost data files](#)

[Accessing Multiple Exchange Mailboxes](#)

[Logging on to Exchange as the Mailbox Owner](#)

[Delegating Mailbox Access](#)

[Opening Additional Exchange Mailboxes](#)

[Granting Permission to Access Folders Without Delegating Access](#)

[Using Mail Profiles to Customize the Mail Environment](#)

[Creating, Copying, and Removing Mail Profiles](#)

[Selecting a Specific Profile to use on Startup](#)

[Chapter 11. Customizing & Troubleshooting the Exchange Shell](#)

[Running and using the Exchange Management Shell](#)

[Managing the PowerShell Application](#)

[Customizing Exchange Management Shell](#)

[Performing One-to-Many Remote Management](#)

[Using a Manual Remote Shell to Work with Exchange](#)

[Preparing to Use the Remote Shell](#)

[Connecting Manually to Exchange 2016 Servers](#)

[Connecting Manually to Exchange Online](#)

[Managing Remote Sessions](#)

[Troubleshooting Exchange Management Shell](#)

[Chapter 12. Customizing & Configuring Exchange Security](#)

[Configuring Standard Exchange Permissions](#)

[Assigning Permissions: Exchange Server and Online](#)

[Understanding Exchange Management Groups](#)

[Assigning Management Permissions](#)

[Understanding Advanced Exchange Server Permissions](#)

[Assigning Advanced Exchange Server Permissions](#)

[Configuring Role-Based Permissions for Exchange](#)

[Understanding Role-Based Permissions](#)

[Working with Role Groups](#)

[Managing Role Group Members](#)

[Assigning Roles Directly or Via Policy](#)

[Configuring Account Management Permissions](#)

[Managing Advanced Permissions](#)

[Adding Custom Roles](#)

[Adding Custom Role Scopes](#)

[Adding Custom Role Entries](#)

[Working with Shared and Split Permissions](#)

[Using Shared Permissions](#)

[Using Split Permissions](#)

[Chapter 13. Implementing Exchange Services](#)

[Selecting Hardware for Exchange 2016](#)

[Navigating Exchange 2016 Editions](#)

[Using Exchange 2016 with Windows Server](#)

[Services for Exchange Server](#)

[Exchange Server Authentication and Security](#)

[Exchange Server Security Groups](#)

[Using Exchange 2016 with Active Directory](#)

[Understanding How Exchange Stores Information](#)

[Understanding How Exchange Routes Messages](#)

[Additional Tools and Options](#)

[Chapter 14. Preparing for Exchange 2016](#)

[Designing the Exchange Server Organization](#)

[Planning for High Availability](#)

[Planning Exchange Databases and Storage](#)

[Planning for Client Access](#)

[Planning to Support Transport Services](#)

[Planning for Unified Messaging](#)

[Integrating Exchange with Active Directory](#)

[How Mailbox Servers use Active Directory](#)

[How Edge Transports use Active Directory](#)

[Integrating Exchange 2016 Into Existing Organizations](#)

[Coexistence and Active Directory](#)

[Configuring Exchange 2016 for Coexistence](#)

[Setting the Default Offline Address Book](#)

[Moving to Exchange Server 2016](#)

[Chapter 15. Deploying Exchange Server 2016](#)

[Installing New Exchange Servers](#)

[Installing Exchange Server](#)

[Verifying and Completing the Installation](#)

[Uninstalling Exchange 2016](#)

[Using Cumulative Updates](#)

[What's in Cumulative Updates?](#)

[How Are Cumulative Updates Applied?](#)

[How Do I Track Exchange Version Numbers?](#)

[Installing Cumulative Updates and Service Packs](#)

[Preparing to Install a Cumulative Update or Service Pack](#)

[Installing a Cumulative Update or Service Pack](#)

[Chapter 16. Exchange 2016 Administration Essentials](#)

[Working with Exchange Admin Center](#)

[Accessing Exchange Admin Center](#)

[Working with Exchange Server Certificates](#)

[Configuring Exchange Admin Center](#)

[Bypassing Exchange Admin Center and Troubleshooting](#)

[Understanding Remote Execution in Exchange Admin Center](#)

[Bypassing Exchange Admin Center and Exchange Management Shell](#)

[Troubleshooting OWA, ECP, Powershell, and More](#)

[Resolving SSL Certificate Issues](#)

[Resolving OWA, ECP, or Other Virtual Directory Issues](#)

[Validating Exchange Server Licensing](#)

[Using and Managing Exchange Services](#)

[Working with Exchange Services](#)

[Checking Required Services](#)

[Maintaining Exchange Services](#)

[Configuring Service Startup](#)

[Configuring Service Recovery](#)

[Customizing Remote Management Services](#)

[Chapter 17. Managing Exchange Organizations](#)

[Navigating Exchange 2016 Organizations](#)

[Organizational Architecture](#)

[Front End Transport](#)

[Back End Transport](#)

[Understanding Exchange Routing](#)

[Routing Boundaries](#)

[IP Site Links](#)

[Cross-Premises Routing](#)

[Understanding Data Storage in Exchange Server 2016](#)

[Working with the Active Directory Data Store](#)

[Using Multimaster Replication](#)

[Using Global Catalogs](#)

[Using Dedicated Expansion Servers](#)

[Navigating the Exchange Information Store](#)

[Data Storage Components](#)

[The Managed Store](#)

[Exchange Server Data Files](#)

[Data Storage in Exchange Databases](#)

[Exchange Server Message Queues](#)

[Chapter 18. Implementing Availability Groups](#)

[Building Blocks for High Availability](#)

[The Extensible Storage Engine](#)

[The High Availability Framework](#)

[Cluster Components](#)

[Active Manager Framework](#)

[Managed Availability Components](#)

[Creating and Managing Database Availability Groups](#)

[Preparing for DAGs](#)

[Creating Database Availability Groups](#)

[Managing Availability Group Membership](#)

[Managing Database Availability Group Networks](#)

[Changing Availability Group Network Settings](#)

[Configuring Database Availability Group Properties](#)

[Removing Servers from a Database Availability Group](#)

[Removing Database Availability Groups](#)

[Maintaining Database Availability Groups](#)

[Switching Over Servers and Databases](#)

[Checking Continuous Replication Status](#)

[Restoring Operations After a DAG Member Failure](#)

[Chapter 19. Configuring Exchange Databases](#)

[Getting Started with Active Mailbox Databases](#)

[Planning for Mailbox Databases](#)

[Preparing for Automatic Reseed](#)

[Creating and Managing Active Databases](#)

[Creating Mailbox Databases](#)

[Setting the Default Offline Address Book](#)

[Setting Mailbox Database Limits and Deletion Retention](#)

[Recovering Deleted Mailboxes](#)

[Recovering Deleted Items from Servers](#)

[Creating and Managing Database Copies](#)

[Creating Mailbox Database Copies](#)

[Configuring Database Copies](#)

[Suspending and Resuming Replication](#)

[Activating Lagged Database Copies](#)

[Updating Mailbox Database Copies](#)

[Monitoring Database Replication Status](#)

[Removing Database Copies](#)

[Maintaining Mailbox Databases](#)

[Checking Database Status](#)

[Setting the Maintenance Interval](#)

[Renaming Databases](#)

[Mounting and Dismounting Databases](#)

[Configuring Automatic Mounting](#)

[Moving Databases](#)

[Deleting Databases](#)

[Managing Content Indexing](#)

[Indexing Essentials](#)

[Maintaining Exchange Store Search](#)

[Resolving Indexing Issues](#)

[Chapter 20. Managing SMTP Connectors](#)

[Send and Receive Connectors: The Essentials](#)

[Understanding Send and Receive Connectors](#)

[Routing Messages within Sites](#)

[Routing Messages Across Site Links](#)

[Managing Send Connectors](#)

[Creating Send Connectors](#)

[Viewing and Managing Send Connectors](#)

[Configuring Send Connector DNS Lookups](#)

[Setting Send Connector Limits](#)

[Managing Receive Connectors](#)

[Creating Receive Connectors](#)

[Configuring Receive Connectors](#)

[Creating Connectors with Exchange Online](#)

[Chapter 21. Configuring Transport Services](#)

[Optimizing Transport Limits](#)

[Setting Organizational Transport Limits](#)

[Setting Connector Transport Limits](#)

[Setting Server Transport Limits](#)

[Setting Exchange Activesync Limits](#)

[Setting Exchange Web Services Limits](#)

[Setting Outlook Web App Limits](#)

[Managing Message Transport](#)

[Configuring the Postmaster Address and Mailbox](#)

[Configuring Shadow Redundancy](#)

[Configuring Safety Net](#)

[Enabling Anti-Spam Features](#)

[Subscribing Edge Transport Servers](#)

[Creating an Edge Subscription](#)

[Getting Edge Subscription Details](#)

[Synchronizing Edge Subscriptions](#)

[Verifying Edge Subscriptions](#)

[Removing Edge Subscriptions](#)

[Chapter 22. Maintaining Mail Flow](#)

[Managing Message Routing and Delivery](#)

[Understanding Message Pickup and Replay](#)

[Configuring and Moving the Pickup and Replay Directories](#)

[Changing the Message Processing Speed](#)

[Configuring Messaging Limits for the Pickup Directory](#)

[Configuring Message Throttling](#)

[Understanding Back Pressure](#)

[Creating and Managing Accepted Domains](#)

[Understanding SMTP Domains](#)

[Viewing Accepted Domains](#)

[Creating Accepted Domains](#)

[Changing The Accepted Domain Type and Identifier](#)

[Removing Accepted Domains](#)

[Creating and Managing Remote Domains](#)

[Viewing Remote Domains](#)

[Creating Remote Domains](#)

[Configuring Messaging Options for Remote Domains](#)

[Removing Remote Domains](#)

[Chapter 23. Implementing Exchange Policies and Rules](#)

[Creating and Managing Email Address Policies](#)

[Working with Email Address Policies](#)

[Creating Email Address Policies](#)

[Editing and Applying Email Address Policies](#)

[Removing Email Address Policies](#)

[Configuring Journal Rules](#)

[Setting The NDR Journaling Mailbox](#)

[Creating Journal Rules](#)

[Managing Journal Rules](#)

[Configuring Transport Rules](#)

[Creating Transport Rules](#)

[Managing Transport Rules](#)

[Chapter 24. Filtering Spam](#)

[Filtering Spam by Sender](#)

[Filtering Spam by Recipient](#)

[Filtering Connections with IP Block Lists](#)

[Applying IP Block Lists](#)

[Configuring Block List Providers](#)

[Specifying Custom Error Messages](#)

[Defining Block Lists](#)

[Using Connection Filter Exceptions](#)

[Using Global Allowed Lists](#)

[Using Global Block Lists](#)

[Preventing Internal Servers from Being Filtered](#)

[Chapter 25. Optimizing Web and Mobile Access](#)

[Navigating IIS Essentials for Exchange Server](#)

[Understanding Mobile Access via IIS](#)

[Maintaining Virtual Directories and Web Applications](#)

[Starting, Stopping, and Restarting Websites](#)

[Configuring Outlook Web App Features](#)

[Managing Segmentation Features](#)

[Managing Outlook Web App Policies](#)

[Managing Bindings, Connections and Authentication](#)

[Optimizing the Mobile Access Websites](#)

[Enabling SSL on Websites](#)

[Restricting Incoming Connections](#)

[Redirecting Users to Alternate Urls](#)

[Controlling Access to the HTTP Server](#)

[Throttling Client Access to Servers](#)

[Optimizing Access for Web and Mobile Clients](#)

[Configuring Access for OAB](#)

[Configuring Access for OWA](#)

[Configuring Access for Exchange ActiveSync](#)

[Configuring Access for ECP](#)

[Chapter 26. Optimizing Client Access Protocols](#)

[Managing RPC and MAPI over HTTP](#)

[Working with RPC and MAPI over HTTP](#)

[Configuring URLs and Authentication](#)

[Enabling the POP3 and IMAP4 Services](#)

[Optimizing POP3 and IMAP4 Settings](#)

[Configuring POP3 and IMAP4 Bindings](#)

[Configuring POP3 and IMAP4 Authentication](#)

[Configuring Connection Settings for POP3 and IMAP4](#)

[Configuring Message Retrieval Settings for POP3 and IMAP4](#)

[Chapter 27. Configuring Mobile Messaging](#)

[Mastering Mobile and Wireless Access Essentials](#)

[Getting Started with Exchange ActiveSync](#)

[Managing ActiveSync and OWA for Devices](#)

[Configuring Autodiscover](#)

[Understanding Autodiscover](#)

[Maintaining Autodiscover](#)

[Using Direct Push](#)

[Using Remote Device Wipe](#)

[Remotely Wiping a Device](#)

[Reviewing the Remote Wipe Status](#)

Using Password Recovery

Recovering a Device Password

Managing File Access and Document Viewing

Configuring Direct File Access

Configuring Remote File Access

Integrating Office Web Apps Servers

Working with Mobile Devices and Device Policies

Viewing Existing Mobile Device Mailbox Policies

Creating Mobile Device Mailbox Policies

Optimizing Mobile Device Mailbox Policies

Assigning Mobile Device Mailbox Policies

Removing Mobile Device Mailbox Policies

Managing Device Access

Blocking Device Access

Using Access Rules

Setting Access Levels and Blocking Thresholds

Chapter 28. Tracking and Logging Exchange Server 2016

Configuring Message Tracking

Changing the Logging Location

Setting Logging Options

Searching the Tracking Logs

Beginning an Automated Search

Reviewing Logs Manually

Searching the Delivery Status Reports

Configuring Protocol Logging

Enabling or Disabling Protocol Logging

Setting Other Protocol Logging Options

Managing Protocol Logging

Optimizing Protocol Logging for HTTP

Working with HTTP Protocol Logs

Using Connectivity Logging

Configuring Connectivity Logging

Working with Connectivity Logs

Chapter 29. Maintaining Exchange Server 2016

Monitoring Events, Services, Servers, and Resource Usage

Viewing Events

Managing Essential Services

Monitoring Messaging Components

Using Performance Alerting

Tracking Memory Usage

Tracking CPU Utilization

Tracking Disk Usage

[Working with Queues](#)

[Understanding Exchange Queues](#)

[Accessing the Queue Viewer](#)

[Managing Queues](#)

[Understanding Queue Summaries and Queue States](#)

[Refreshing The Queue View](#)

[Working with Messages In Queues](#)

[Forcing Connections to Queues](#)

[Suspending and Resuming Queues](#)

[Deleting Messages from Queues](#)

[**Chapter 30. Troubleshooting Exchange Server 2016**](#)

[Troubleshooting Essentials](#)

[Tracking Server Health](#)

[Tracking User and Workload Throttling](#)

[Tracking Configuration Changes](#)

[Testing Service Health, Mail Flow, Replication and More](#)

[Diagnosing and Resolving Problems](#)

[Identifying Recovery Actions](#)

[Identifying Responders](#)

[Identifying Monitors](#)

[Identifying Probes](#)

[Viewing Error Messages for Probes](#)

[Tracing Probe Errors](#)

[Troubleshooting Outlook Web App](#)

[Checking OWA Health](#)

[Understanding Unhealthy Status](#)

[Correcting Unhealthy Status](#)

[**Index**](#)

[**About the Author**](#)

How to Use This Guide

This guide combines the complete text of *Exchange Server 2016 & Exchange Online: Essentials for Administration, 2nd Edition* and *Exchange Server 2016: Server Infrastructure, 2nd Edition*, providing 30 chapters and 250,000 words of in-depth insights into Microsoft's latest enterprise messaging server. The first chapter, *Welcome to Exchange Server 2016*, delivers the essential context for working with the product. Chapters detailing how to put Exchange Server 2016 to work follow.

William Stanek has been developing expert solutions for and writing professionally about Microsoft Exchange since 1995. In this book, William shares his extensive knowledge of the product, delivering ready answers for day-to-day management and zeroing in on core commands and techniques.

This book is written especially for IT Professionals working with, supporting and managing Exchange Server 2016, Exchange Online and Office 365. Because Exchange Online and Office 365 are online products, the features and options for these products can be updated from time to time by Microsoft.

Print Readers

Print editions of this book include an index and some other elements not available in the digital edition. Updates to this book are available online. Visit <http://www.williamrstanek.com/exchangeserver/> to get any updates. This content is available to all readers.

Digital Book Readers

Digital editions of this book are available at all major retailers, at libraries upon request and with many subscription services. If you have a digital edition of this book that you downloaded elsewhere, such as a file sharing site, you should know that the author doesn't receive any royalties or income from such downloads. Already downloaded this book or others? Donate here to ensure William can keep writing the books you need:

https://www.paypal.com/cgi-bin/webscr?cmd=_s-xclick&hosted_button_id=CPSBGLZ35AB26

Support Information

Every effort has been made to ensure the accuracy of the contents of this book. As corrections are received or changes are made, they will be added to the online page for the book available at:

<http://www.williamrstanek.com/exchangeserver/>

If you have comments, questions, or ideas regarding the book, or questions that are not

answered by visiting the site above, send them via e-mail to:

williamstanek@aol.com

Other ways to reach the author:

Facebook: <http://www.facebook.com/William.Stanek.Author>

Twitter: <http://twitter.com/williamstanek>

It's important to keep in mind that Microsoft software product support is not offered. If you have questions about Microsoft software or need product support, please contact Microsoft.

Microsoft also offers software product support through the Microsoft Knowledge Base at:

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

Conventions & Features

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

The first letters of the names of menus, dialog boxes, user interface elements, and commands are capitalized. Example: the New Mail Contact dialog box. This book also has notes, tips and other sidebar elements that provide additional details on points that need emphasis.

Keep in mind that throughout this book, where William has used click, right-click and double-click, you can also use touch equivalents, tap, press and hold, and double tap. Also, when using a device without a physical keyboard, you are able to enter text by using the onscreen keyboard. If a device has no physical keyboard, simply touch an input area on the screen to display the onscreen keyboard.

Share & Stay in Touch

The marketplace for technology books has changed substantially over the past few years. In addition to becoming increasingly specialized and segmented, the market has been shrinking rapidly, making it extremely difficult for books to find success. If you want William to be able to continue writing and write the books you need for your career, raise your voice and support his work.

Without support from you, the reader, future books by William will not be possible. Your voice matters. If you found the book to be useful, informative or otherwise helpful, please take the time to let others know by sharing about the book online.

To stay in touch with William, visit him on Facebook or follow him on Twitter. William welcomes messages and comments about the book, especially suggestions for

improvements and additions. If there is a topic you think should be covered in the book, let William know.

Chapter 1. Welcome to Exchange 2016

Before getting to the specifics of working with Exchange 2016, take a few moments to familiarize yourself with the configuration options available. Microsoft Exchange is available in on-premises, online and hybrid implementations.

With an on-premises implementation, you deploy Exchange server hardware on your network and manage all aspects of the implementation. Here, you control the servers and determine which version of Exchange those servers will run. Exchange Server 2016 is the current version of Exchange, and was released in its original implementation in October 2015. Like other releases of Exchange, Exchange Server 2016 is updated periodically with software updates which may change or enhance the options available.

With an online implementation, you manage the service-level settings, organization configuration, and recipient configuration while relying on Microsoft for hardware and other services. Microsoft determines which version of Exchange those servers will run. Online implementations always use the most current release version of Exchange, which at present is Exchange Server 2016. As with any on-premises implementation, Microsoft's servers are updated periodically with software updates which may change or enhance the options available.

Although either an on-premises or online implementation can be your only solution for all your enterprise messaging needs, a hybrid implementation gives you an integrated online and on-premises solution. Here, your organization controls the on-premises servers and Microsoft controls the online servers. The on-premises servers and online servers can run the same, or different, versions of Exchange.

Understanding the various implementation scenarios will help you work through the rest of this book and will also help you navigate Exchange 2016 and its management options on your own. This chapter covers the basics. You'll learn about Exchange Admin Center and Exchange Management Shell, the essential tools for managing Exchange 2016.

As you get started with Exchange 2016, it's important to point out that with this version, Microsoft has completed the consolidation of server roles begun with Exchange 2013. Thus, Mailbox and Edge Transport are now the only server roles available. Mailbox servers now perform all messaging and client access tasks except for perimeter security, which can be handled by servers running the Edge Transport role.

Getting Started with Exchange Admin Center

Exchange Admin Center replaces Exchange Management Console and Exchange Control Panel (ECP) used in early releases of Exchange and there is no longer a separate graphical administration tool. Exchange Admin Center is a browser-based application designed for managing on-premises, online, and hybrid Exchange organizations.

For on-premises management, you access Exchange Admin Center through the Mailbox servers deployed in your Exchange organization. For online management, you access Exchange Admin Center through the Mailboxes servers hosted by Microsoft. Although the application can be configured with an internal access URL and a separate external access URL, only an internal access URL is configured by default in on-premises configurations. This means that by default you can access Exchange Admin Center only when you are on the corporate network.

Navigating Exchange Admin Center Options

After you log in to Exchange Admin Center, you'll see the list view with manageable features listed in the left pane, also called the Navigation menu (see Figure 1-1). When you select a feature in the Navigation menu, you'll then see the related topics or "tabs" for that feature. The manageable items for a selected topic or tab are displayed in the main area of the browser window. For example, when you select Recipients in the Navigation menu, the topics or tabs that you can work with are: Mailboxes, Groups, Resources, Contacts, Shared and Migration.

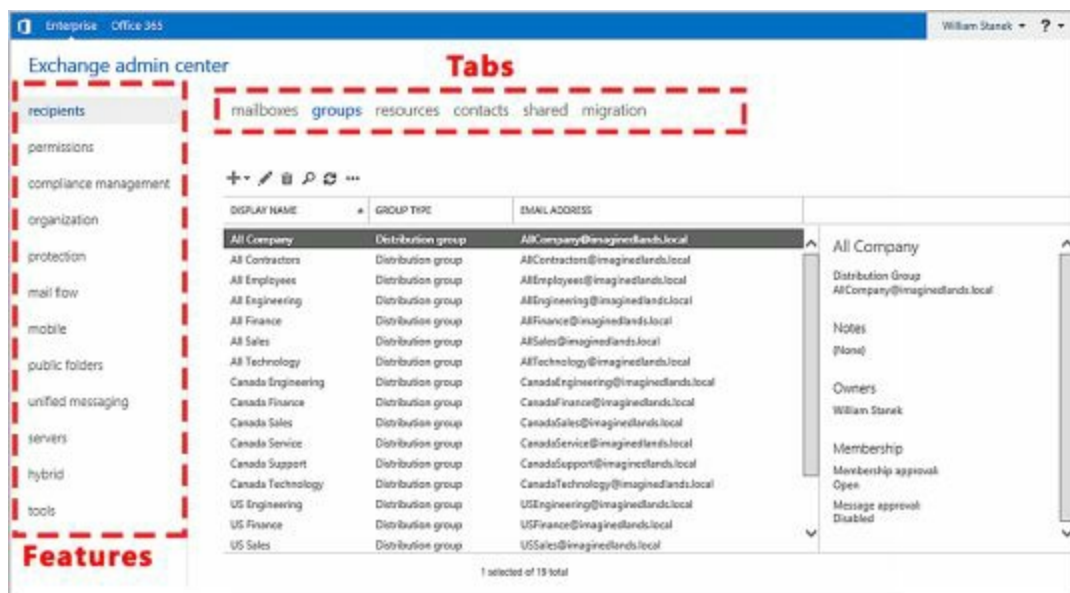



FIGURE 1-1 Exchange Admin Center features and tabs

Below the tabs, you'll find a row of Option buttons:

 **New** – Allows you to create a new item.



Edit – Allows you to edit a selected item.



Delete – Deletes a selected item.



Search – Performs a search within the current context.



Refresh – Refreshes the display so you can see changes.



More – If available, displays additional options.

As shown in Figure 1-2, the navigation bar at the top of the window has several important options. You use the Enterprise and Office 365 options for cross-premises navigation.

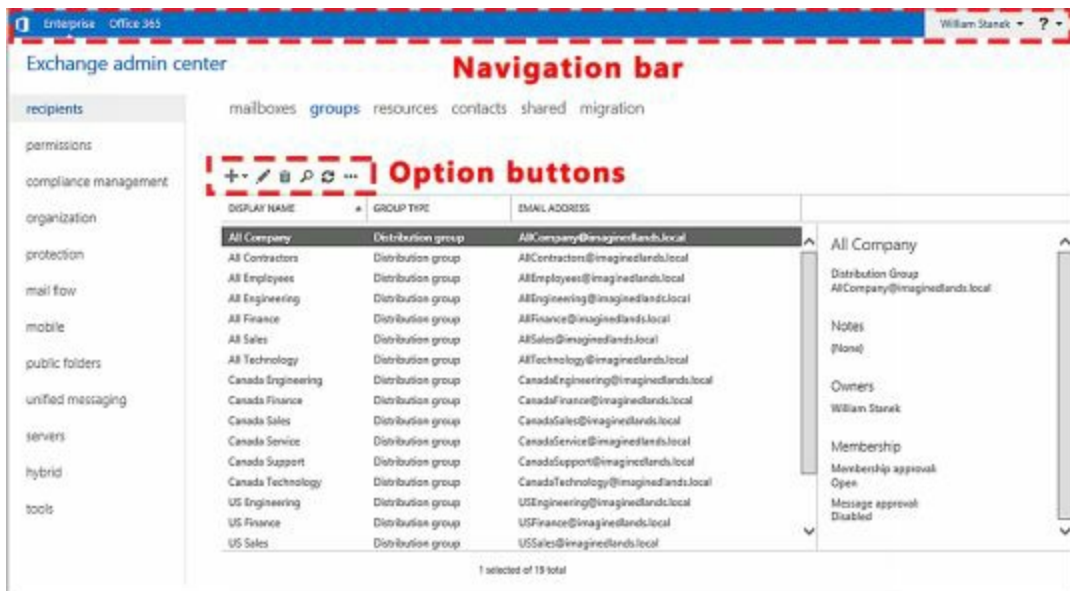


FIGURE 1-2 The Navigation bar in Exchange Admin Center

If there are notifications, you'll see a Notification icon on the Navigation bar. Clicking this icon displays notifications, such as alerts regarding automated or batch processes. The User button shows the currently logged on user. Clicking the User button allows you to logout or sign in as another user.

As shown in Figure 1-3, when working with recipients, such as mailboxes or groups, you can click the More button (**⋮**) to display options to:

- Add or remove columns
- Export data for the listed recipients to a .csv file
- Perform advanced searches

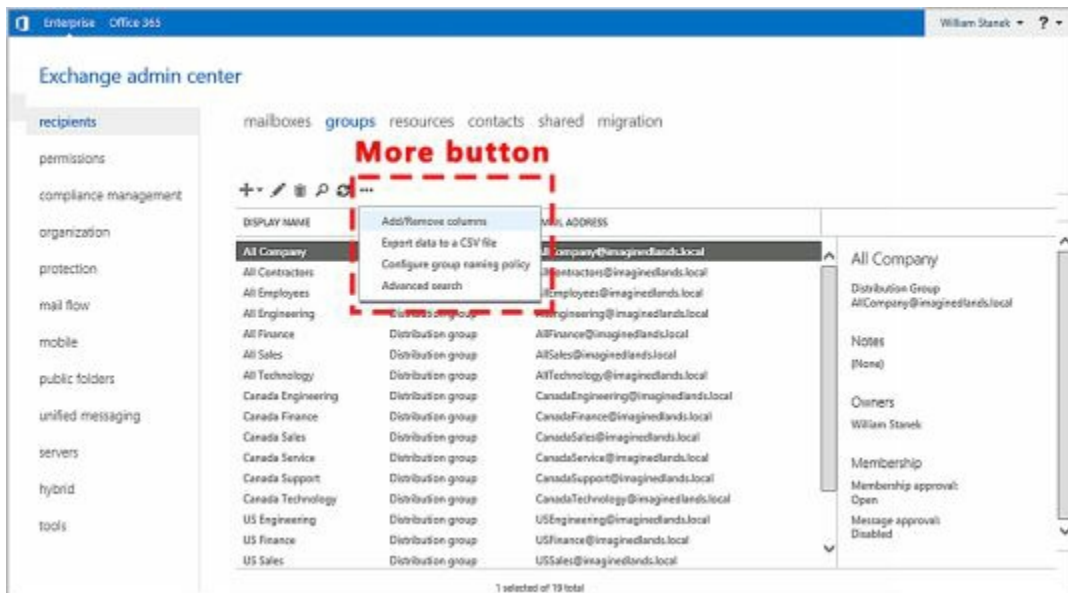


FIGURE 1-3 The More button in Exchange Admin Center

If you customize the view by adding or removing columns, the settings are saved for the computer that you are using to access Exchange Admin Center. However, because the settings are saved as browser cookies, clearing the browser history will remove the custom settings.

When working with recipients, you typically can select multiple items and perform bulk editing as long as you select like items, such as mailbox users or mail-enabled contacts. Select multiple items using the Shift or Ctrl key and then use bulk editing options in the Details pane to bulk edit the selected items.

NOTE Although ECP for Exchange 2010 would return only 500 recipients at a time, Exchange Admin Center for Exchange 2016 doesn't have this limitation. Results are paged so that you can go through results one page at a time and up to 20,000 recipients can be returned in the result set.

Accessing Exchange Admin Center

Exchange Admin Center is designed to be used with Windows, Windows Server and other operating systems. When you are working with Windows or Windows Server, you can use Internet Explorer or the Edge browser. With other operating systems, such as Linux, you can use Firefox or Chrome. On Mac OS X 10.5 or later, you can also use Safari.

You access Exchange Admin Center by following these steps:

1. Open your web browser and enter the secure URL for Exchange Admin Center. If you are outside the corporate network, enter the external URL, such as <https://mail.imagedlands.com/ecp>. If you are inside the corporate network, enter the internal URL, such as <https://mailserver23/ecp>.
2. If your browser displays a security alert stating there's a problem with the site's security certificate or that the connection is untrusted, proceed anyway. This alert is displayed because the browser does not trust the self-signed certificate that

was created when the Exchange server was installed.

- With Internet Explorer, the error typically states "There's a problem with this website's security certificate." Proceed by selecting the Continue To This Web Site (Not Recommended) link.
 - With Google Chrome, the error typically states "The site's security certificate is not trusted." Continue by clicking Proceed Anyway.
 - With Mozilla Firefox, the error typically states "This connection is untrusted." Proceed by selecting I Understand The Risks and then selecting Add Exception. Finally, in the Add Security Exception dialog box, select Confirm Security Exception.
3. You'll see the logon page for Exchange Admin Center (as shown in Figure 1-4). Enter your user name and password, then click **Sign In**.

Be sure to specify your user name in DOMAIN\username format. The domain can either be the DNS domain, such as imaginedlands.com, or the NetBIOS domain name, such as pocket-consulta. For example, the user AnneW could specify her logon name as imaginedlands.com\annew or imaginedlands\annew.

4. If you are logging on for the first time, select your preferred display language and time zone, and then click **Save**.



FIGURE 1-4 Signing in to Exchange Admin Center

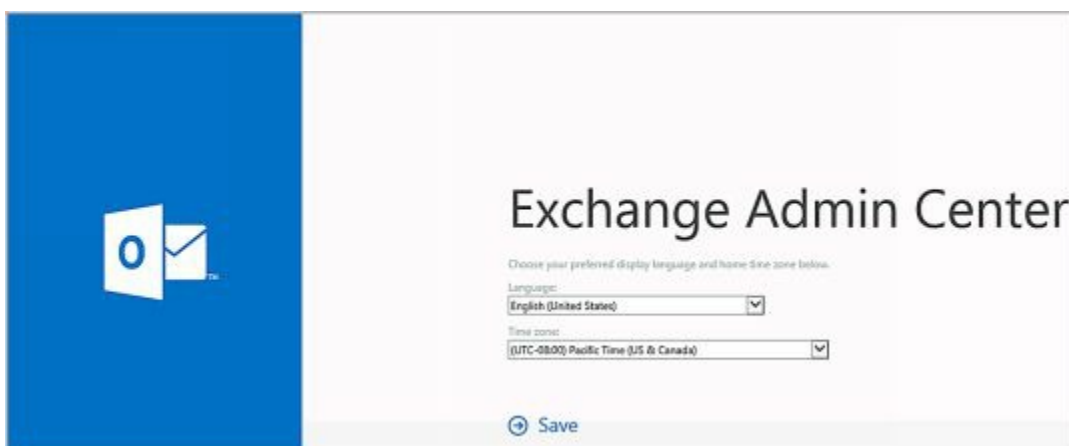


FIGURE 1-5 Setting the language and time zone.

To ensure all features are available you should only use Exchange Admin Center with the most recent version of the browser available for your operating system. The version of Exchange Admin Center you see depends on the version of Exchange running on the

Mailbox server hosting your personal mailbox. Exchange 2016 runs version 15.1, and you can specify this version explicitly by appending **?ExchClientVer=15** to the internal or external URL. By default, you must use HTTPS to connect. Using HTTPS ensures data transmitted between the client browser and the server is encrypted and secured.

Authenticating and Proxying Connections

When you access Exchange Admin Center in a browser, a lot is happening in the background that you don't see. Although you access the application using a specific server in your organization, the Client Access service running on the server acts as a front-end proxy that authenticates and proxies the connection to the Exchange back end using Internet Information Services (IIS). Thus, although Mailbox server functions perform the actual back-end processing, the front-end IIS configuration is essential to proper operations.

As shown in Figure 1-6, you can examine the configuration settings for Exchange Admin Center and other applications using Internet Information Services (IIS) Manager. The server to which you connect processes your remote actions via the ECP application running on the default website. The physical directory for this application is `%ExchangeInstallPath%\FrontEnd\HttpProxy\Ecp`. This application runs in the context of an application pool named `MSExchangeECPAppPool`. In the `%ExchangeInstallPath%\FrontEnd\HttpProxy\Ecp` directory on your server, you'll find a `web.config` file that defines the settings for the ECP application.

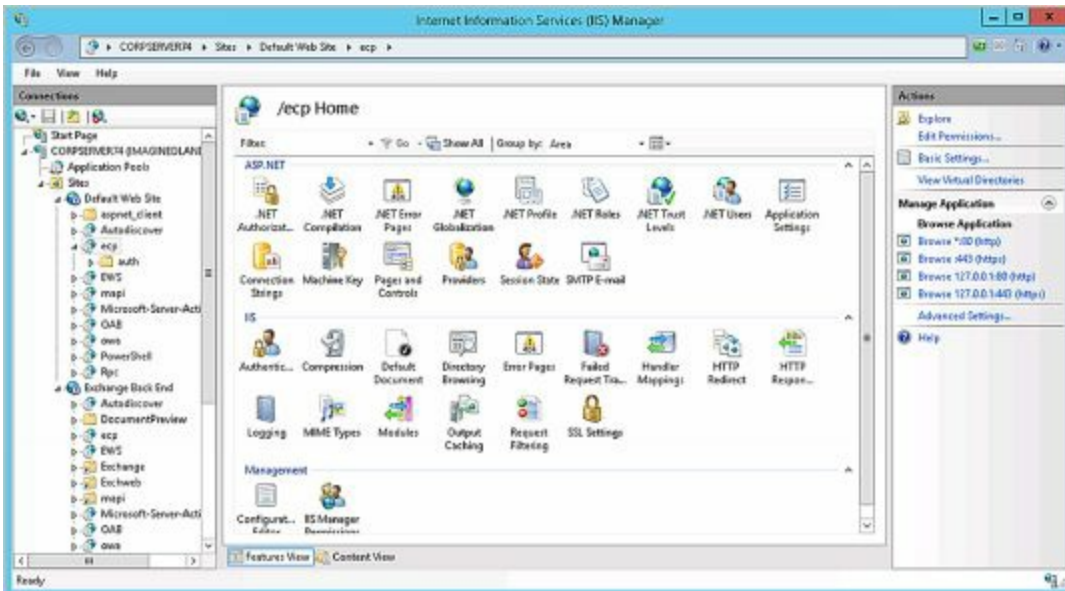


FIGURE 1-6 Viewing the applications that handle Exchange processing.

The Mailbox server where your mailbox resides performs its tasks and processing via the ECP application running on the Exchange Back End website. The physical directory for this application is `%ExchangeInstallPath%\ClientAccess\Ecp`. This application also runs in the context of an application pool named `MSExchangeECPAppPool`. In the `%ExchangeInstallPath%\ClientAccess\Ecp` directory on your server, you'll find a `web.config` file that defines the settings for the ECP application.

Getting Started with Exchange Management Shell

Microsoft Exchange Server 2016 includes Exchange Management Shell, which is an extensible command-line environment for Exchange Server that builds on the existing framework provided by Windows PowerShell. When you install Exchange Server 2016 on a server, or when you install the Exchange Server management tools on a workstation, you install Exchange Management Shell as part of the process.

When you start Exchange Management Shell, the working environment is loaded automatically with many of the working environment features coming from profiles, which are a type of script that runs automatically when you start the shell. However, the working environment is also determined by other imported elements.

Running and Using Cmdlets

A cmdlet(pronounced *commandlet*) is the smallest unit of functionality when working with command shells. You can think of a cmdlet as a built-in command. Rather than being highly complex, most cmdlets are quite simple and have a small set of associated properties.

You use cmdlets the same way you use any other commands and utilities. Cmdlet names are not case sensitive. This means you can use a combination of both uppercase and lowercase characters. After starting the shell, you can type the name of the cmdlet at the prompt, and it will run in much the same way as a command-line command.

For ease of reference, cmdlets are named using verb-noun pairs. The verb tells you what the cmdlet does in general. The noun tells you what specifically the cmdlet works with. For example, the Get-Variable cmdlet gets a named environment variable and returns its value. If you don't specify which variable to get as a parameter, Get-Variable returns a list of all environment variables and their values.

You can work with cmdlets in several ways:

- [Executing commands directly at the shell prompt](#)
- [Running commands from scripts](#)
- [Calling them from C# or other .NET Framework languages](#)

You can enter any command or cmdlet you can run at the shell prompt into a script by copying the related command text to a file and saving the file with the .ps1 extension. You can then run the script in the same way you would any other command or cmdlet. Keep in mind that when you are working with Windows PowerShell, the current directory is not part of the environment path in most instances. Because of this, you typically need to use “./” when you run a script in the current directory, such as:

```
./runtasks
```

NOTE Windows PowerShell includes a rich scripting language and allows the use of standard language constructs for looping, conditional execution, flow

control, and variable assignment. Discussion of these features is beyond the scope of this book. A good resource is *Windows PowerShell: The Personal Trainer* .

Running and Using Other Commands and Utilities

Because the shell runs within the context of the Windows command prompt, you can run all Windows command-line commands, utilities, and graphical applications from within the shell. However, remember that the shell interpreter parses all commands before passing off the command to the command prompt environment. If the shell has a like-named command or a like-named alias for a command, this command, and not the expected Windows command, is executed.

Non-shell commands and programs must reside in a directory that is part of the PATH environment variable. If the item is found in the path, it is run. The PATH variable also controls where the shell looks for applications, utilities, and scripts. In the shell, you can work with Windows environment variables using `$env`. To view the current settings for the PATH environment variable, type `$env:path` . To add a directory to this variable, use the following syntax:

```
$env:path += "; DirectoryPathToAdd "
```

where *DirectoryPathToAdd* is the directory path you want to add to the path, such as:

```
$env:path += ";C:\Scripts"
```

To have this directory added to the path every time you start the shell, you can add the command line as an entry in your profile. Profiles store frequently used elements, including aliases and functions. Generally speaking, profiles are always loaded when you work with the shell. Keep in mind that cmdlets are like built-in commands rather than standalone executables. Because of this, they are not affected by the PATH environment variable.

Using Cmdlet Parameters and Errors

You use parameters to control the way cmdlets work. All cmdlet parameters are designated with an initial dash (–). As some parameters are position-sensitive, you sometimes can pass parameters in a specific order without having to specify the parameter name. For example, with `Get-Service`, you don't have to specify the `–Name` parameter and can simply type:

```
get-service ServiceName
```

where *ServiceName* is the name of the service you want to examine, such as:

```
get-service MSExchangeIS
```

Here, the command returns the status of the Microsoft Exchange Information Store service. Because you can use wildcards, such as `*`, with name values, you can also type `get-service mse*` to return the status of all Microsoft Exchange-related services.

When you work with cmdlets, you'll encounter two standard types of errors: terminating

errors and nonterminating errors. While terminating errors halt execution, nonterminating errors cause error output to be returned but do not halt execution. With either type of error, you'll typically see error text that can help you resolve the problem that caused it. For example, an expected file might be missing or you might not have sufficient permissions to perform a specified task.

Using Cmdlet Aliases

For ease of use, the shell lets you create aliases for cmdlets. An alias is an abbreviation for a cmdlet that acts as a shortcut for executing the cmdlet. For example, you can use the alias `gsv` instead of the cmdlet name `Get-Service`.

At the shell prompt, enter **get-alias** to list all currently defined aliases. Define additional aliases using the `Set-Alias` cmdlet. The syntax is:

```
set-alias aliasName cmdletName
```

where *aliasName* is the alias you want to use and *cmdletName* is the cmdlet for which you are creating an alias. The following example creates a “go” alias for the `Get-Process` cmdlet:

```
set-alias go get-process
```

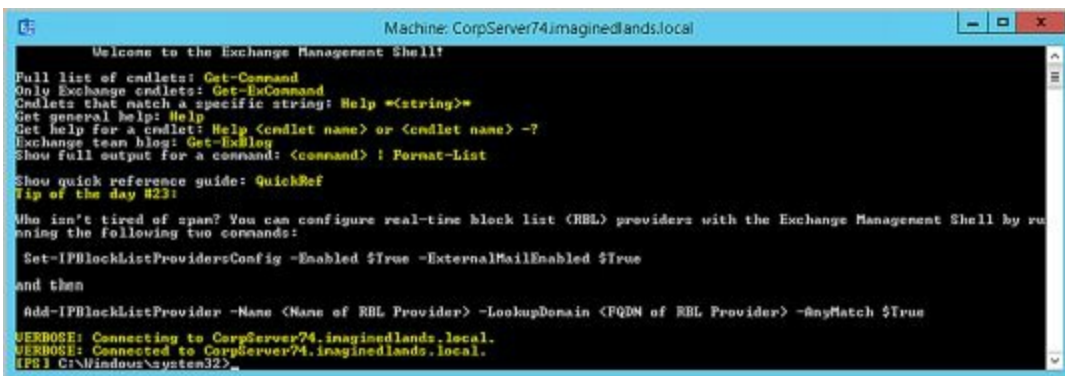
To use your custom aliases whenever you work with the shell, enter the related command line in your profile.

Working with Exchange Management Shell

The Exchange Management Shell is a command-line management interface built on Windows PowerShell. You use the Exchange Management Shell to manage any aspect of an Exchange Server 2016 configuration that you can manage in the Exchange Admin Center. This means that you can typically use either tool to configure Exchange Server 2016. However, only the Exchange Management Shell has the full complement of available commands, and this means that some tasks can be performed only at the shell prompt.

Starting Exchange Management Shell

After you've installed the Exchange management tools on a computer, the Exchange Management Shell, shown in Figure 1-7, is available. On desktop computers running Windows 8.1 or Windows 10, one way to start the shell is by using the Apps Search box. Type **shell** in the Apps Search box, and then select Exchange Management Shell. Or click Start, click All Apps and then choose Exchange Management Shell.



```
Machine: CorpServer74\imaginedlands.local
Welcome to the Exchange Management Shell!
Full list of cmdlets: Get-Command
Only Exchange cmdlets: Get-ExCommand
Cmdlets that match a specific string: Help <string>
Get general help: Help
Get help for a cmdlet: Help <cmdlet name> or <cmdlet name> -?
Exchange team blog: Get-ExBlog
Show full output for a command: <command> | Format-List
Show quick reference guide: QuickRef
Tip of the day #23:
Who isn't tired of spam? You can configure real-time block list (RBL) providers with the Exchange Management Shell by running the following two commands:
Set-IPBlockListProvidersConfig -Enabled $True -ExternalMailEnabled $True
and then
Add-IPBlockListProvider -Name <Name of RBL Provider> -LookupDomain <FQDN of RBL Provider> -AnyMatch $True
VERBOSE: Connecting to CorpServer74.imaginedlands.local.
VERBOSE: Connected to CorpServer74.imaginedlands.local.
PS C:\Windows\system32>
```

FIGURE 1-7 Exchange Management Shell

Starting Exchange Management Shell on Windows Server 2012 R2 or Windows Server 2016 is a little different. Here, click Start and then click the More button (**...**) to display the Apps screen. On the Apps screen, choose Exchange Management Shell. While you are working with the Apps screen, right-click the tile for Exchange Management Shell and then select Pin To Start or Pin To Taskbar. This will make it easier to access the shell in the future.

Exchange Management Shell is designed to be run only on domain-joined computers. Whether you are logged on locally to an Exchange server or working remotely, starting the shell opens a custom Windows PowerShell console. The console does the following:

1. Connects to the closest Exchange 2016 server using Windows Remote Management (WinRM).
2. Performs authentication checks that validate your access to the Exchange 2016 server and determine the Exchange role groups and roles your account is a member of. You must be a member of at least one management role.

3. Creates a remote session with the Exchange 2016 server. A remote session is a runspace that establishes a common working environment for executing commands on remote computers.

NOTE It's important to note that selecting the shell in this way starts the Exchange Management Shell using your user credentials. This enables you to perform any administrative tasks allowed for your user account and in accordance with the Exchange role groups and management roles you're assigned. As a result, you don't need to run the Exchange Management Shell in administrator mode, but you can. To do so, right-click Exchange Management Shell shortcut, and then click Run As Administrator.

Using Exchange Cmdlets

When you are working with the Exchange Management Shell, additional Exchange-specific cmdlets are available. As with Windows PowerShell cmdlets, you can get help information on Exchange cmdlets:

- To view a list of all Exchange cmdlets, enter **get-excommand** at the shell prompt.
- To view a list of Exchange cmdlets for a particular item, such as a user, contact or mailbox, enter **get-help *ItemName***, where *ItemName* is the name of the item you want to examine.

When you work with the Exchange Management Shell, you'll often work with Get, Set, Enable, Disable, New, and Remove cmdlets (the groups of cmdlets that begin with these verbs). These cmdlets all accept the `-Identity` parameter, which identifies the unique object with which you are working.

Typically, a cmdlet that accepts the `-Identity` parameter has this parameter as its first parameter, allowing you to specify the identity, with or without the parameter name. When identities have names as well as aliases, you can specify either value as the identity. For example, you can use any of the following techniques to retrieve the mailbox object for the user William Stanek with the mail alias Williams:

```
get-mailbox -identity william
get-mailbox -identity 'William Stanek'
get-mailbox Williams
get-mailbox "William Stanek"
```

With Get cmdlets, you typically can return an object set containing all related items simply by omitting the identity. For example, if you type **get-mailbox** at the shell prompt without specifying an identity, you get a list of all mailboxes in the enterprise (up to the maximum permitted to return in a single object set).

By default, all cmdlets return data in table format. Because there are often many more columns of data than fit across the screen, you might need to switch to Format-List output to see all of the data. To change to the Format-List output, redirect the output using the pipe symbol (`|`) to the Format-List cmdlet, as shown in this example:

```
get-mailbox -identity williams | format-list
```

You can abbreviate Format-List as *fl* , as in this example:

```
get-mailbox -identity williams | fl
```

Either technique typically ensures that you see much more information about the object or the result set than if you were retrieving table-formatted data.

Working with Object Sets and Redirecting Output

When you are working with PowerShell or Exchange Management Shell, you'll often need to redirect the output of one cmdlet and pass it as input to another cmdlet. You can do this using the pipe symbol. For example, if you want to view mailboxes for a specific mailbox database rather than all mailboxes in the enterprise, you can pipe the output of Get-MailboxDatabase to Get-Mailbox, as shown in this example:

```
get-mailboxdatabase -Identity "Engineering" | get-mailbox
```

Here, you use Get-MailboxDatabase to get the mailbox database object for the Engineering database. You then send this object to the Get-Mailbox cmdlet as input, and Get-Mailbox iterates through all the mailboxes in this database. If you don't perform any other manipulation, the mailboxes for this database are listed as output, as shown here:

Name	Alias	Server	ProhibitSendQuota
Administrator	Administrator	mailboxsvr82	unlimited
William S	williams	mailboxsvr82	unlimited
Tom G	tomg	mailboxsvr82	unlimited
David W	davidw	mailboxsvr82	unlimited
Kari F	karif	mailboxsvr82	unlimited
Connie V	conniev	mailboxsvr82	unlimited
Mike D	miked	mailboxsvr82	unlimited

You can also pipe this output to another cmdlet to perform an action on each individual mailbox in this database. If you don't know the name of the mailbox database you want to work with, enter **get-mailboxdatabase** without any parameters to list all available mailbox databases.

Chapter 2. Working with Exchange Online

Exchange Online is available as part of an Office 365 plan and as a standalone service. Microsoft offers a variety of Office 365 plans that include access to Office Web Apps, the full desktop versions of Office, or both as well as access to Exchange Online. You'll likely want to use an Office 365 midsize business or enterprise plan to ensure Active Directory integration is included as you'll need this feature to create a hybrid Exchange organization. If you don't want to use Office 365, Microsoft offers plans specifically for Exchange Online. The basic plans are the cheapest but don't include in-place hold and data loss prevention features that large enterprises may need to meet compliance and regulatory requirements. That said, both basic and advanced plans support Active Directory integration for synchronization with on-premises Active Directory infrastructure and the creation of hybrid Exchange organizations.

In Exchange Online, email addresses, distribution groups, and other directory resources are stored in the directory database provided by Active Directory for Windows Azure. Windows Azure is Microsoft's cloud-based server operating system. Exchange Online fully supports the Windows security model and by default relies on this security mechanism to control access to directory resources. Because of this, you can control access to mailboxes and membership in distribution groups and perform other security administration tasks through the standard permission set.

Because Exchange Online uses Windows security, you can't create a mailbox without first creating a user account that will use the mailbox. Every Exchange mailbox must be associated with a user account—even those used by Exchange Online for general messaging tasks.

Getting Started with Exchange Online

With Exchange Online, the tools you'll use most often for administration are Office Admin Center and Exchange Admin Center. Regardless of whether you use Exchange Online with Office 365, you'll use Office Admin Center as it's where you manage service-level settings, including the Office tenant domain, subscriptions, and licenses.

Navigating Exchange Online Services

When you sign up for Office 365 and Exchange Online, you'll be provided an access URL, such as <https://portal.microsoftonline.com/admin/default.aspx>. After you log in by entering your username and password, you'll see the Office Admin Center dashboard, shown in Figure 2-1.

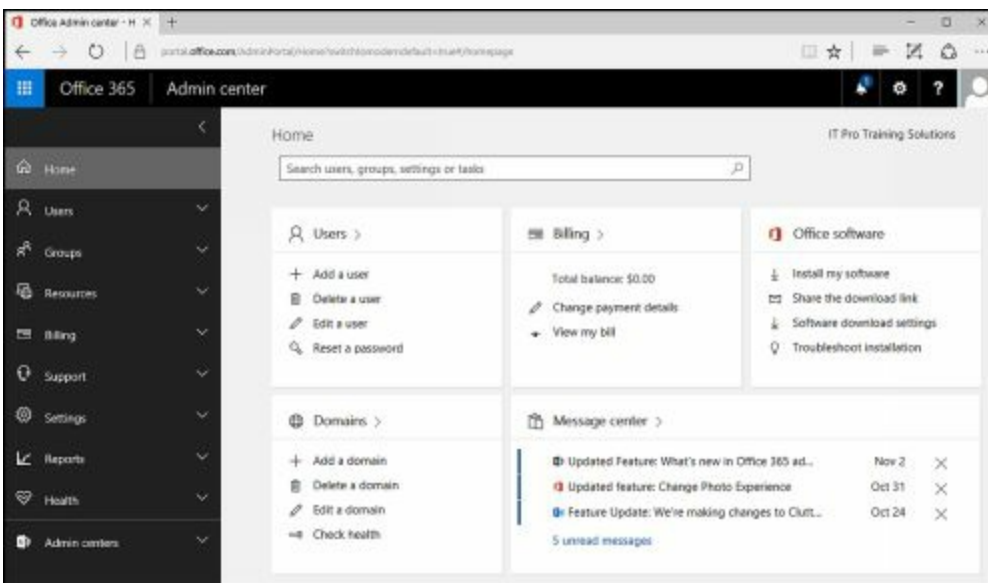


FIGURE 2-1 Use Office Admin Center to manage users and accounts

From the Office Admin Center dashboard, you have full access to Office 365 and Exchange Online.

As with Exchange Admin Center, Office Admin Center has a horizontal Navigation bar with several options:




Apps – Displays a list of the available apps you can switch to, including Office Admin Center.



Notifications – Displays notifications, such as alerts regarding licensing or subscription issues.



Help – Displays help and feedback options.



Settings – Displays options for accessing account settings.



Account – Displays the name of the currently logged in user and provides options for accessing the account's profile page and signing out.

IMPORTANT As you get started with Exchange Online, it's important to keep in mind that available features and options can change over time. Why? Microsoft releases cumulative updates for Exchange on a fixed schedule and applies these cumulative updates to their hosted Exchange servers prior to official release of an update for on-premises Exchange servers. Thus, when you see that an update has been released for the current Exchange Server product you know it has been applied to all Exchange Online servers and all of the mailboxes stored in the cloud as well.

Below the horizontal Navigation bar, you'll find the Navigation menu, shown in Figure 2-2.

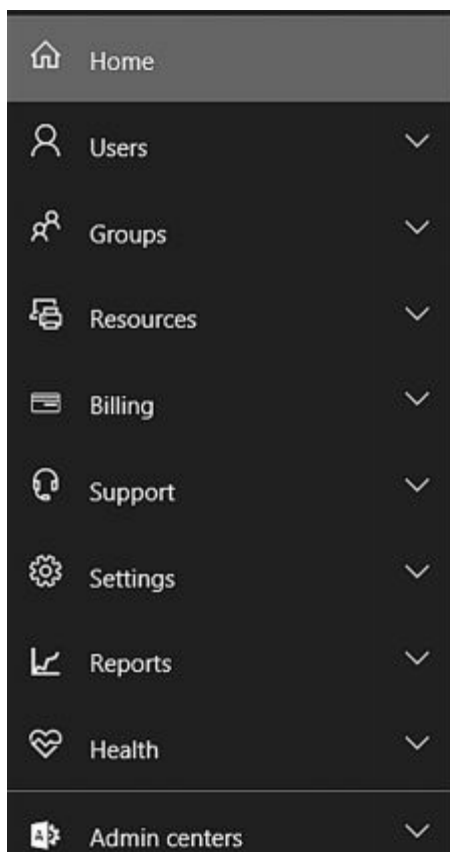


FIGURE 2-2 The Navigation menu in Office 365 Admin Center

Just as the Navigation menu can be expanded by clicking:



Or collapsed by clicking:




Each item on the menu can be expanded by clicking:



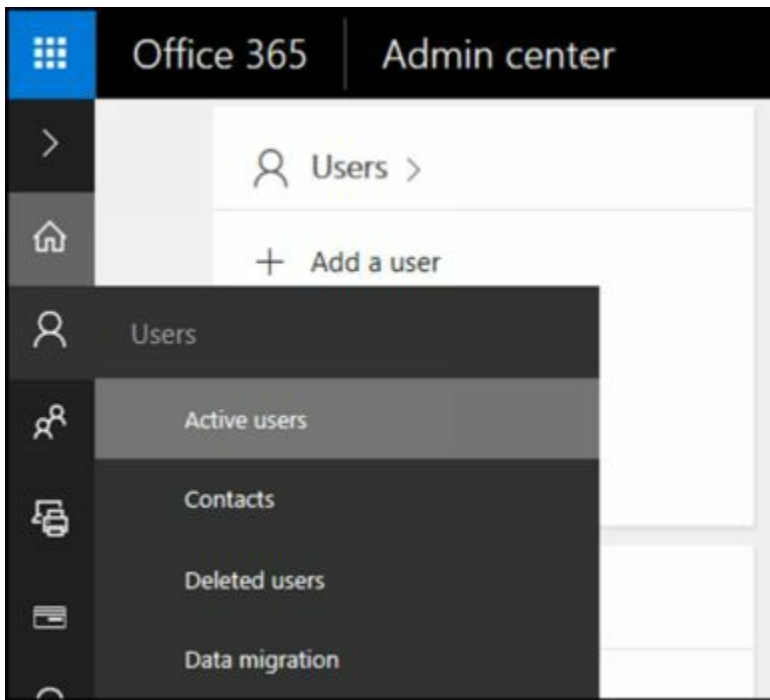
Or collapsed by clicking:



When you expand a menu item, you see a list of related options.



The same options are available when you hover over or point to a menu item with the Navigation menu collapsed.




Like Office Admin Center, Exchange Admin Center for Exchange Online is a web application. You use Exchange Admin Center for Exchange Online to manage:

- **Organization configuration data.** This type of data is used to manage policies, address lists, and other types of organizational configuration details.
- **Recipient configuration data.** This type of data is associated with mailboxes, mail-enabled contacts, and distribution groups.

Although Exchange Admin Center for on-premises installations and Exchange Admin Center for Exchange Online are used in the same way and have many similarities, they also have many differences. These differences include limitations that apply to the online environment but do not apply to on-premises environments.

The easiest way to access Exchange Admin Center for Exchange Online, shown in Figure 2-3, is via Office Admin Center:



1. If the Admin Centers () panel is closed, click or hover over it in the Navigation menu to see the related options.
2. Click Exchange under the Admin Centers heading. This opens the Exchange Admin Center dashboard.

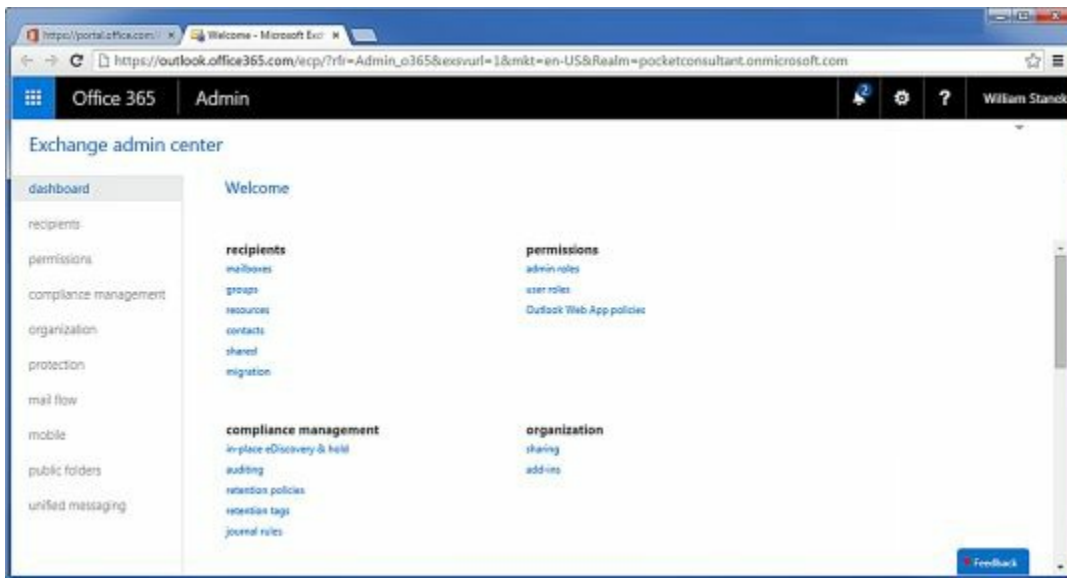


FIGURE 2-3 Use the Exchange Admin Center to manage recipients, permissions and more.

If you are using a mixed or hybrid environment with on-premises and online services, you can also access Office Admin Center and Exchange Admin Center for Exchange Online from your on-premises installation:

1. Access Exchange Admin Center for your on-premises installation and then click Office 365 on the Navigation bar.
2. After your browser connects to Office.com, click Sign In on the Navigation bar and then select Work, School Or University as your account type.
3. Provide the email address and password for your Microsoft account and then click Sign In. This opens Office Admin Center.
4. In Office Admin Center, click Menu and then click Exchange under the Admin Centers heading. This opens the Exchange Admin Center dashboard.

The dashboard is unique to Exchange Admin Center for Exchange Online and serves to provide quick access to commonly used features. These features are also available via the Navigation menu and the related tabs.

Other than the dashboard, Exchange Admin Center for Exchange Online works just like Exchange Admin Center for on-premises installations. Manageable features are listed in the Navigation menu. After you select a feature in the Navigation menu, you'll see the related topics or "tabs" for that feature. The manageable items for a selected topic or tab are displayed in the main area of the browser window. For example, when you select Organization in the Navigation menu, the topics or tabs that you can work with are: Sharing and Apps.


Understanding Office 365 Licensing

With Exchange Online, you perform administration using either Exchange Admin Center or Windows PowerShell—not Exchange Management Shell, which is meant to be used only with on-premises installations of Exchange. Regardless of which approach you use to create new users in Exchange Online, you must license mailbox users in Office 365.

You do this by licensing mailbox plans and associating a mailbox plan with each mailbox user.

Using Exchange Admin Center, you can associate mailbox plans when you create mailbox users or afterward by editing the account properties. In PowerShell, you use the `New-Mailbox` cmdlet with the `-MailboxPlan` parameter to do the same.

When you assign mailbox plans, you need to ensure you have enough licenses. You purchase and assign licenses using Office 365 Admin Center:

1. If the Billing options aren't currently displayed in the Navigation menu, expand Billing () by clicking or hovering over it in the Navigation menu and then click Licensing to see the number of valid, expired and assigned licenses.
2. Click Subscriptions under Billing in the Navigation menu to display subscription and licensing options.
3. Click Add Subscriptions to purchase additional services. For example, if you scroll down the list of purchasable services, you'll see the Exchange Online plans.
4. While viewing plans, click Buy Now to purchase a particular plan. As shown in Figure 2-4, you'll have the option to specify how many user licenses you want for the selected plan before you check out.

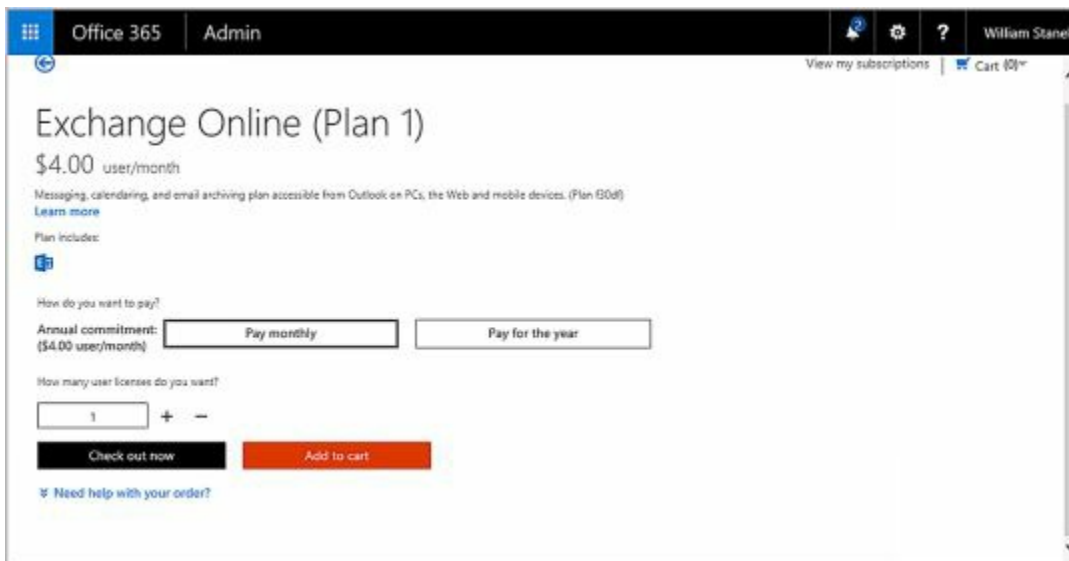


FIGURE 2-4 Select a plan and purchase licenses.

Although Office 365 will allow you to assign more mailbox plans than you have licenses for, you shouldn't do this. After the initial grace period, problems will occur. For example, mail data for unlicensed mailboxes may become unavailable. Remember, the number of valid licenses shouldn't exceed the number of assigned licenses.

You activate and license synced users in Office 365 as well. Under `Users > Active Users`, select the check boxes for the users you want to activate and license and then select `Activate Synced Users`. Next, specify the work location for the users, such as `United States`. Under `Assign Licenses`, select the mailbox plan to assign. Finally, select `Activate`.

Using Windows PowerShell with Exchange Online

Although Office Admin Center and Exchange Admin Center provide everything you need to work with Exchange Online, there may be times when you want to work from the command line, especially if you want to automate tasks with scripts. Enter Windows PowerShell.

Getting Started with Windows PowerShell

Windows PowerShell is built into Windows and Windows Server. Windows PowerShell supports cmdlets, functions and aliases. Cmdlets are built-in commands. Functions provide basic functionality. Aliases are abbreviations for cmdlet names. As cmdlet, function and alias names are not case sensitive, you can use a combination of both uppercase and lowercase characters to specify cmdlet, function and alias names.

Although Windows PowerShell has a graphical environment called Windows PowerShell ISE (`powershell_ise.exe`), you'll usually work with the command-line environment. The PowerShell console (`powershell.exe`) is available as a 32-bit or 64-bit environment for working with PowerShell at the command line. On 32-bit versions of Windows, you'll find the 32-bit executable in the `%SystemRoot%\System32\WindowsPowerShell\v1.0` directory.

On 64-bit versions of Windows and Windows Server, a 64-bit and a 32-bit console are available. The default console is the 64-bit console, which is located in the `%SystemRoot%\System32\WindowsPowerShell\v1.0` directory. The 32-bit executable in the `%SystemRoot%\SysWow64\WindowsPowerShell\v1.0` directory and is labeled as Windows PowerShell (x86).

With Windows 8.1 or later, you can start the PowerShell console by using the Apps Search box. Type **powershell** in the Apps Search box, and then press Enter. Or you can select Start and then choose Windows PowerShell. From Mac OS X or Linux, you can run either Windows 7 or later in a virtual environment to work with Windows PowerShell.

In Windows, you also can start Windows PowerShell from a command prompt (`cmd.exe`) by typing **powershell** and pressing Enter. To exit Windows PowerShell and return to the command prompt, type `exit`.

When the shell starts, you usually will see a message similar to the following:

```
Windows PowerShell
Copyright (C) 2016 Microsoft Corporation.
All rights reserved.
```

You can disable this message by starting the shell with the `-Nologo` parameter, such as:

```
powershell -nologo
```

By default, the version of scripting engine that starts depends on the operating system

you are using. With Windows 8.1 and Windows Server 2012 R2, the default scripting engine is version 4.0. With Windows 10 and Windows Server 2016, the default scripting engine is version 4.0. To confirm the version of Windows PowerShell installed, enter the following command:

```
Get-Host | Format-List Version
```

Because you can abbreviate Format-List as FL, you also could enter:

```
Get-Host | fl Version
```

NOTE Letter case does not matter with Windows PowerShell. Thus, Get-Host, GET-HOST and get-host are all interpreted the same.

Figure 2-5 shows the PowerShell window. When you start PowerShell, you can set the version of the scripting engine that should be loaded. To do this, use the `-Version` parameter. In this example, you specify that you want to use PowerShell Version 3.0:

```
powershell -version 3
```

NOTE Windows can only load available versions of the scripting engine. For example, you won't be able to run the version 5.1 scripting engine if only PowerShell Version 4.0 is available.



FIGURE 2-5 Use the PowerShell console to manage Exchange remotely at the prompt.

By default, the PowerShell window displays 50 lines of text and is 120 characters wide. When additional text is to be displayed in the window or you enter commands and the PowerShell console's window is full, the current text is displayed in the window and prior text is scrolled up. To temporarily pause the display when a command is writing output, press `Ctrl+S`. You can then press `Ctrl+S` to resume or `Ctrl+C` to terminate execution.

Understanding the Default Working Environment

When you run Windows PowerShell, a default working environment is loaded automatically. The features for this working environment come primarily from profiles, which are a type of script that run automatically whenever you start PowerShell. The working environment also is determined by imported snap-ins, providers, modules, command paths, file associations, and file extensions.

To start Windows PowerShell without loading profiles, use the `-NoProfile` parameter,

such as:

```
powershell -noprofile
```

Whenever you work with scripts, you need to keep in mind the current execution policy and whether signed scripts are required. Execution policy is a built-in security feature of Windows PowerShell that controls whether and how you can run configuration files and scripts. Although the default configuration depends on which operating system and edition are installed, policy is always set on either a per-user or per-computer basis in the Windows registry.

You can display the execution policy currently being applied, using the `Get-ExecutionPolicy` cmdlet. The available execution policies, from least secure to most secure, are:

- **Bypass.** Bypasses warnings and prompts when scripts run. Use with programs that have their own security model or when a PowerShell script is built into a larger application.
- **Unrestricted.** Allows all configuration files and scripts to run whether they are from local or remote sources and regardless of whether they are signed or unsigned. When you run a configuration file or script from a remote resource, you are prompted with a warning that the file comes from a remote resource before the configuration file is loaded or the script runs.
- **RemoteSigned.** Requires all configuration files and scripts from remote sources to be signed by a trusted publisher. However, configuration files and scripts on the local computer do not need to be signed. PowerShell does not prompt you with a warning before running scripts from trusted publishers.
- **AllSigned.** Requires all configuration files and scripts from all sources—whether local or remote—to be signed by a trusted publisher. Thus, configuration files and scripts on the local computer and remote computers must be signed. PowerShell prompts you with a warning before running scripts from trusted publishers.
- **Restricted.** Prevents PowerShell from loading configuration files and scripts. Effects all configuration files and scripts, regardless of whether they are signed or unsigned. Because a profile is a type of script, profiles are not loaded either.
- **Undefined.** Removes the execution policy that is set for the current user scope and instead applies the execution policy set in Group Policy or for the LocalMachine scope. If execution policy in all scopes is set to Undefined, the default execution policy, Restricted, is the effective policy.

By default, when you set execution policy, you are using the LocalMachine scope, which is applied to all users of the computer. You also can set the scope to CurrentUser so that the execution policy level is only applied to the currently logged on user.

Using `Set-ExecutionPolicy`, you can change the preference for the execution policy. Normally, changes to execution policy are written to the registry. However, if the Turn On Script Execution setting in Group Policy is enabled for the computer or user, the user preference is written to the registry, but it is not effective. Windows PowerShell will

display a message explaining that there is a conflict. Finally, you cannot use Set-ExecutionPolicy to override a group policy, even if the user preference is more restrictive than the policy setting. For example, you can set the execution policy to run scripts regardless of whether they have a digital signature and work in an unrestricted environment by entering:

```
set-executionpolicy unrestricted
```

When you change execution policy, the change occurs immediately and is applied to the local console or application session. Because the change is written to the registry, the new execution policy normally will be used whenever you work with PowerShell.

Learning About Cmdlets and Functions

When you are working with Windows PowerShell, you can get a complete list of cmdlets and functions available by entering **get-command**. The output lists cmdlets and functions by name and associated module.

Another way to get information about cmdlets is to use Get-Help. When you enter **get-help *-***, you get a list of all cmdlets, including a synopsis that summarizes the purpose of the cmdlet. Rather than listing help information for all commands, you can get help for specific commands by following Get-Help with the name of the cmdlet you want to work with, such as:

```
get-help clear-history
```

Because Windows PowerShell V3 and later use online and updatable help files, you may see only basic syntax for cmdlets and functions when you use Get-Help. To get full help details, you'll have to either use online help or download the help files to your computer. For online help, add the `-online` parameter to your Get-Help command, such as:

```
get-help get-variable -online
```

You can use the Update-Help cmdlet to download and install the current help files from the Internet. Without parameters, Update-Help updates the help files for all modules installed on the computer. When you are working with Update-Help, keep the following in mind:

- Update-Help downloads files only once a day
- Update-Help only installs help files when they are newer than the ones on the computer
- Update-Help limits the total size of uncompressed help files to 1 GB

You can override these restrictions using the `-Force` parameter.

Connecting to Exchange Online Using PowerShell

The way you use Windows PowerShell to manage Exchange Server and Exchange Online are different. With Exchange Server installations, you manage Exchange using Exchange Management Shell, which is a command-line management interface built on Windows PowerShell that you can use to manage any aspect of an Exchange Server configuration that you can manage in the Exchange Admin Center. With Exchange Online installations, you manage Exchange using a remote session and the built-in functions and capabilities of Exchange Management Shell are not available.

Exploring How the Shell Uses Remote Sessions

The Exchange Management Shell is designed to be run only on domain-joined computers and is available when you have installed the Exchange management tools on a management computer or server. Whether you are logged on locally to an Exchange server or working remotely, starting Exchange Management Shell opens a custom Windows PowerShell console that runs in a remote session with an Exchange server.

A remote session is a runspace that establishes a common working environment for executing commands on remote computers. Before creating the remote session, this custom console connects to the closest Exchange server using Windows Remote Management (WinRM) and then performs authentication checks that validate your access to the Exchange server and determine the Exchange role groups and roles your account is a member of. You must be a member of at least one management role.

Because the Exchange Management Shell uses your user credentials, you are able to perform any administrative tasks allowed for your user account and in accordance with the Exchange role groups and management roles you're assigned. You don't need to run the Exchange Management Shell in elevated, administrator mode, but you can by right-clicking Exchange Management Shell, and then selecting Run As Administrator.

By examining the properties of the shortcut that starts the Exchange Management Shell, you can see the actual command that runs when you start the shell is:

```
C:\Windows\System32\WindowsPowerShell\v1.0\powershell.exe -noexit -command  
". 'C:\Program Files\Microsoft\Exchange Server\V15\bin\RemoteExchange.ps1';  
Connect-ExchangeServer -auto -ClientApplication:ManagementShell "
```

Here, the command starts PowerShell, runs the RemoteExchange.ps1 profile file, and then uses the command Connect-ExchangeServer to establish the remote session. The –Auto parameter tells the cmdlet to automatically discover and try to connect to an appropriate Exchange server. The –ClientApplication parameter specifies that client-side application is the Exchange Management Shell. When you run the shell in this way, Windows Powershell loads a profile script called RemoteExchange.ps1 that sets aliases, initializes Exchange global variables, and loads .NET assemblies for Exchange. The profile script also modifies the standard PowerShell prompt so that it is scoped to

the entire Active Directory forest and defines Exchange-specific functions, including:

- **Get-Exbanner.** Displays the Exchange Management Shell startup banner.
- **Get-Exblog.** Opens Internet Explorer and accesses the Exchange blog.
- **Get-Excommand.** Lists all available Exchange commands.
- **Get-Pscommand.** Lists all available PowerShell commands.
- **Get-Tip.** Displays the tip of the day.
- **Quickref.** Opens Internet Explorer and accesses the Exchange Management Shell quick start guide.

All of these processes simplify the task of establishing an interactive remote session with Exchange server. As implemented in the default configuration, you have a one-to-one, interactive approach for remote management, meaning you establish a session with a specific remote server and work with that specific server whenever you execute commands.

Establishing Remote Sessions

When you are working with PowerShell outside of Exchange Management Shell, you must manually establish a remote session with Exchange. As the RemoteExchange.ps1 profile file and related scripts are not loaded, the related cmdlets and functions are not available. This means you cannot use Get-Exbanner, Get-Exblog, Get-Excommand, Get-Pscommand, Get-Tip or Quickref. Further, when you are working with an online installation of Exchange, the cmdlets available are different from when you are working with Exchange Server.

PowerShell provides several cmdlets for establishing remote sessions, including Enter-PSSession and New-PSSession. The difference between the two options is subtle but important.

Using an Interactive Remote Session

You can use the Enter-PSSession cmdlet to start an interactive session with Exchange or any other remote computer. The basic syntax is Enter-PSSession ComputerName, where ComputerName is the name of the remote computer, such as the following:

```
enter-pssession Server58
```

When the session is established, the command prompt changes to show that you are connected to the remote computer, as shown in the following example:

```
[Server58]: PS C:\Users\wrstaneck.cpand\Documents>
```

While working in a remote session, any commands you enter run on the remote computer just as if you had typed them directly on the remote computer. Generally, to perform administration, you need to use an elevated, administrator shell and pass credentials along in the session. Establishing a connection in this way uses the standard PowerShell remoting configuration.

However, you cannot connect to Exchange Online using the standard PowerShell remoting configuration. You must go through a PowerShell application running on ps.outlook.com or another appropriate web server. Typically, when you work with Exchange Online, you use the connection URI `https://ps.outlook.com/powershell/` and the actual session is redirected to your specific online server. To ensure redirection doesn't fail, you must add the `-AllowRedirection` parameter.

As shown in the following example, you use the `-ConnectionUri` parameter to specify the connection URI, the `-ConfigurationName` parameter to specify the configuration namespace, and the `-Authentication` parameter to set the authentication type to use:

```
Enter-PSSession -ConfigurationName Microsoft.Exchange  
-ConnectionUri https://ps.outlook.com/powershell/  
-Authentication Basic -AllowRedirection
```

Here, you set the configuration namespace as `Microsoft.Exchange`, establish a connection to the Exchange Online URL provided by Microsoft, and use Basic authentication. As you don't specify credentials, you will be prompted to provide credentials.

You also can pass in credentials as shown in this example:

```
Enter-PSSession -ConfigurationName Microsoft.Exchange  
-ConnectionUri https://ps.outlook.com/powershell/  
-Authentication Basic -Credential  
wrstanek@imaginedlands.onmicrosoft.com  
-AllowRedirection
```

Here, you pass in credentials and are prompted for the associated password.

Alternatively, you can store credentials in a Credential object and then use `Get-Credential` to prompt for the required credentials, as shown here:

```
$Cred = Get-Credential  
Enter-PSSession -ConfigurationName Microsoft.Exchange  
-ConnectionUri https://ps.outlook.com/powershell/  
-Authentication Basic -Credential  
$Cred -AllowRedirection
```

When you are finished working with Exchange Online, you can end the interactive session by using `Exit-PSSession` or by typing `exit`. Although `Enter-PSSession` provides a quick and easy way to establish a remote session, the session ends when you use `Exit-PSSession` or `exit` the PowerShell prompt and there is no way to reestablish the original session. Thus, any commands you are running and any command context is lost when you exit the session.

Thus, as discussed in this section, the basic steps for using a standard interactive remote session are:

1. Open an administrator Windows PowerShell prompt.
2. Use `Enter-PSSession` to establish a remote session.

3. Work with Exchange Online.
4. Exit the remote session using `Exit-PSSession` or by exiting the PowerShell window.

Creating and Importing a Remote Session

Instead of using a standard interactive session, you may want to create a session that you disconnect and reconnect. To do this, you establish the session using `New-PSSession` and then import the session using `Import-PSSession`. The basic syntax:

```
$Session = New-PSSession -ConfigurationName  
Microsoft.Exchange -ConnectionUri  
https://ps.outlook.com/powershell/  
-Authentication Basic -Credential  
wrs@imaginedlands.onmicrosoft.com  
-AllowRedirection
```

In this example, you use `New-PSSession` to create a session and store the related object in a variable called `$Session`. You create the session by setting the configuration namespace as `Microsoft.Exchange`, establishing a connection to the Exchange Online URL provided by Microsoft, which typically is `https://ps.outlook.com`, and using HTTPS with Basic authentication for the session. You also allow redirection. Allowing redirection is important as otherwise the session will fail when the Microsoft web servers redirect the session to the actual location of your Exchange Online installation.

To establish the connection, you must always pass in your Exchange Online user name and password. In the previous example, you specify the user name to use and are prompted for the related password. You also could specify the credentials explicitly, as shown here:

```
$Cred = Get-Credential  
$Session = New-PSSession -ConfigurationName Microsoft.Exchange  
-ConnectionUri https://ps.outlook.com/powershell/  
-Authentication Basic -Credential $Cred  
-AllowRedirection
```

Here, you store credentials in a `Credential` object and then use `Get-Credential` to prompt for the required credentials.

After you establish a session with Exchange Online, you must import the server-side PowerShell session into your client-side session. To do this, you enter the following command:

```
Import-PSSession $Session
```

Where `$Session` is the name of the variable in which the session object is stored. You can then work with the remote server and Exchange Online.

When you are finished working remotely, you should disconnect the remote shell. It's important to note that, beginning with Windows PowerShell 3.0, sessions are persistent by default. When you disconnect from a session, any command or scripts that are running

in the session continue running, and you can later reconnect to the session to pick up where you left off. You also can reconnect to a session if you were disconnected unintentionally, such as by a temporary network outage.

Exchange Online allows each administrative account to have up to three simultaneous connections to sever-side sessions. If you close the PowerShell window without disconnecting from the session, the connection remains open for 15 minutes and then disconnects automatically.

To disconnect a session manually without stopping commands or releasing resources, you can use `Disconnect-PSSession`, as shown in this example:

```
Disconnect-PSSession $Session
```

Here, the `$Session` object was instantiated when you created the session and you disconnect while the session continues to be active. As long as you don't exit the PowerShell window in which this object was created, you can use this object to reconnect to the session by entering:

```
Connect-PSSession $Session
```

Later, when you are finished working with Exchange Online, you should remove the session. Removing a session stops any commands or scripts that are running, ends the session, and releases the resources the session was using. You can remove a session by running the following command:

```
Remove-PSSession $Session
```

Thus, as discussed in this section, the basic steps for working with an imported session are:

1. Open an administrator Windows PowerShell prompt.
2. Use `New-PSSession` to establish the remote session.
3. Import the session using `Import-PSSession`.
4. Work with Exchange Online. Optionally, disconnect from the session using `Disconnect-PSSession` and reconnect to the session using `Connect-PSSession`.
5. Remove the remote session using `Remove-PSSession`.

Connecting to Windows Azure

You can manage the Office 365 service, its settings and accounts using either Office Admin Center or Windows PowerShell. Every account you create in the online environment is in fact created in the online framework within which Office 365 and Exchange Online operate. This framework is called Windows Azure, and like Windows Server, it uses directory services provided by Active Directory.

Before you can manage Office 365, its settings, and accounts from Windows PowerShell, you must install the Windows Azure Active Directory module (the current version of which is available at the Microsoft Download Center:

<http://go.microsoft.com/fwlink/p/?linkid=236297>). Any computer capable of running Exchange or acting as a management computer can run this module. However, there are several prerequisites, including .NET Framework 3.5 and the Microsoft Online Services Sign-in Assistant. At the time of this writing, the sign-in assistant was available at <https://www.microsoft.com/en-us/download/details.aspx?id=41950>. Be sure to download and install only the 64-bit versions of the module and the sign-in assistant.

NOTE Although there are newer versions of the .NET Framework available, version 1.0.8 (the current implementation as of this writing) of the Windows Azure Active Directory module requires version 3.5. Install .NET Framework 3.5 as a feature, even if your computer has a newer version of the .NET Framework installed.

After you download and install the required components, the Windows Azure Active Directory module is available for your use in any PowerShell window. This module also is referred to as the Microsoft Online module. Although Windows PowerShell 3.0 and later implicitly import modules, you may need to explicitly import this module in some configurations. After you import the module, if necessary, you can connect to the Windows Azure and Microsoft Online Services using the `Connect-MSOLService` cmdlet.

Because you'll typically want to store your credentials in a Credential object rather than be prompted for them, the complete procedure to connect to Microsoft Online Services by using Windows PowerShell is:

```
import-module msonline
$cred = get-credential
connect-msolservice -credential:$cred
```

Or, with Windows PowerShell 3.0 or later, use:

```
$cred = get-credential
connect-msolservice -credential:$cred
```

After connecting to the service, you can use cmdlets for Windows Azure Active

Directory to manage online settings and objects. For example, if you want to get a list of user accounts that have been created in the online service along with their licensing status, enter **get-msoluser**. The results will be similar to the following:

UserPrincipalName	DisplayName	isLicensed
wrstanek@imaginedlands.onm...	William Stanek	True
tonyv@imaginedlands.onm...	Tony Vidal	False

Cmdlets for Windows Azure Active Directory

Exchange Online runs on Windows Azure rather than Windows Server. As the two operating environments have different directory services, you must use cmdlets specific to Active Directory for Windows Azure if you want to work with users, groups and related objects.

You'll find complete information about these cmdlets online at <http://msdn.microsoft.com/library/azure/jj151815.aspx>. The available cmdlets include:

• Cmdlets for managing groups and roles

- Add-MsolGroupMember
- Add-MsolRoleMember
- Get-MsolGroup
- Get-MsolGroupMember
- Get-MsolRole
- Get-MsolRoleMember
- Get-MsolUserRole
- New-MsolGroup
- Redo-MsolProvisionGroup
- Remove-MsolGroup
- Remove-MsolGroupMember
- Remove-MsolRoleMember
- Set-MsolGroup

• Cmdlets for managing licenses and subscriptions

- Get-MsolAccountSku
- Get-MsolSubscription
- New-MsolLicenseOptions
- Set-MsolUserLicense

• Cmdlets for managing service principals

- Get-MsolServicePrincipal
- Get-MsolServicePrincipalCredential
- New-MsolServicePrincipal
- New-MsolServicePrincipalAddresses

New-MsolServicePrincipalCredential
Remove-MsolServicePrincipal
Remove-MsolServicePrincipalCredential
Set-MsolServicePrincipal

• **Cmdlets for managing users**

Convert-MsolFederatedUser
Get-MsolUser
New-MsolUser
Redo-MsolProvisionUser
Remove-MsolUser
Restore-MsolUser
Set-MsolUser
Set-MsolUserPassword
Set-MsolUserPrincipalName

• **Cmdlets for managing the Azure service**

Add-MsolForeignGroupToRole
Connect-MsolService
Get-MsolCompanyInformation
Get-MsolContact
Get-MsolPartnerContract
Get-MsolPartnerInformation
Redo-MsolProvisionContact
Remove-MsolContact
Set-MsolCompanyContactInformation
Set-MsolCompanySettings
Set-MsolDirSyncEnabled
Set-MsolPartnerInformation

• **Cmdlets for managing domains**

Confirm-MsolDomain
Get-MsolDomain
Get-MsolDomainVerificationDns
Get-MsolPasswordPolicy
New-MsolDomain
Remove-MsolDomain
Set-MsolDomain
Set-MsolDomainAuthentication
Set-MsolPasswordPolicy

- **Cmdlets for managing single sign-on**

Convert-MsolDomainToFederated

Convert-MsolDomainToStandard

Get-MsolDomainFederationSettings

Get-MsolFederationProperty

New-MsolFederatedDomain

Remove-MsolFederatedDomain

Set-MsolADFSContext

Set-MsolDomainFederationSettings

Update-MsolFederatedDomain

You also can enter **get-help *msol*** to get a list of commands specific to Microsoft Online Services.

Working with Exchange Online Cmdlets

When you work with Exchange Online, the operating environment is different from when you are working with on-premises Exchange Server installations. As a result, different cmdlets and options are available.

Cmdlets Specific to Exchange Online

Because the operating environment for Exchange Online is different from on-premises Exchange, Exchange Online has cmdlets that aren't available when you are working with on-premises Exchange. You'll find complete information about these cmdlets online at [https://technet.microsoft.com/library/jj200780\(v=exchg.160\).aspx](https://technet.microsoft.com/library/jj200780(v=exchg.160).aspx). The additional cmdlets include:

• Cmdlets for working with online recipients

- Add-RecipientPermission
- Get-LinkedUser
- Get-RecipientPermission
- Get-RemovedMailbox
- Get-SendAddress
- Import-ContactList
- Remove-RecipientPermission
- Set-LinkedUser
- Undo-SoftDeletedMailbox

• Cmdlets for working with connected accounts

- Get-ConnectSubscription
- Get-HotmailSubscription
- Get-ImapSubscription
- Get-PopSubscription
- Get-Subscription
- New-ConnectSubscription
- New-HotmailSubscription
- New-ImapSubscription
- New-PopSubscription
- New-Subscription
- Remove-ConnectSubscription
- Remove-Subscription
- Set-ConnectSubscription
- Set-HotmailSubscription
- Set-ImapSubscription
- Set-PopSubscription

- **Cmdlets for working with antispam and anti-malware**

Disable-HostedContentFilterRule
Enable-HostedContentFilterRule
Get-HostedConnectionFilterPolicy
Get-HostedContentFilterPolicy
Get-HostedContentFilterRule
Get-HostedOutboundSpamFilterPolicy
Get-QuarantineMessage
New-HostedConnectionFilterPolicy
New-HostedContentFilterPolicy
New-HostedContentFilterRule
Release-QuarantineMessage
Remove-HostedConnectionFilterPolicy
Remove-HostedContentFilterPolicy
Remove-HostedContentFilterRule
Set-HostedConnectionFilterPolicy
Set-HostedContentFilterPolicy
Set-HostedContentFilterRule
Set-HostedOutboundSpamFilterPolicy

- **Cmdlets for working with connectors**

Get-InboundConnector
Get-OutboundConnector
New-InboundConnector
New-OutboundConnector
Remove-InboundConnector
Remove-OutboundConnector
Set-InboundConnector
Set-OutboundConnector

- **Cmdlets for working with messaging policy and compliance**

Get-DataClassificationConfig
Get-RMSTrustedPublishingDomain
Import-RMSTrustedPublishingDomain
Remove-RMSTrustedPublishingDomain
Set-RMSTrustedPublishingDomain

- **Cmdlets for organization and perimeter control**

Enable-OrganizationCustomization
Get-PerimeterConfig
Set-PerimeterConfig

• Cmdlets for online reporting

Get-ConnectionByClientTypeDetailReport
Get-ConnectionByClientTypeReport
Get-CsActiveUserReport
Get-CsAVConferenceTimeReport
Get-CsConferenceReport
Get-CsP2PAVTimeReport
Get-CsP2PSessionReport
Get-GroupActivityReport
Get-MailboxActivityReport
Get-MailboxUsageDetailReport
Get-MailboxUsageReport
Get-MailDetailDlpPolicyReport
Get-MailDetailMalwareReport
Get-MailDetailSpamReport
Get-MailDetailTransportRuleReport
Get-MailFilterListReport
Get-MailTrafficPolicyReport
Get-MailTrafficReport
Get-MailTrafficSummaryReport
Get-MailTrafficTopReport
Get-MessageTrace
Get-MessageTraceDetail
Get-MxRecordReport
Get-OutboundConnectorReport
Get-RecipientStatisticsReport
Get-ServiceDeliveryReport
Get-StaleMailboxDetailReport
Get-StaleMailboxReport

Although cmdlets specific to Windows Azure Active Directory and Exchange Online itself are available, many of the cmdlets associated with on-premises Exchange continue to be available as well. Primarily, these cmdlets include those that are specific to recipients and mailboxes and do not include those specific to Exchange on-premises configurations or to Exchange server configurations. For example, you can continue to use cmdlets for working with mailboxes, including Disable-Mailbox, Enable-Mailbox, Get-Mailbox, New-Mailbox, Remove-Mailbox, and Set-Mailbox. However, you cannot use cmdlets for working with mailbox databases. In Exchange Online, mailbox databases are managed automatically as part of the service.

Working with Exchange Online Cmdlets

When you work with the Exchange Online, you'll often use Get, Set, Enable, Disable,

New, and Remove cmdlets. The groups of cmdlets that begin with these verbs all accept the `-Identity` parameter, which identifies the unique object with which you are working. Generally, these cmdlets have the `-Identity` parameter as the first parameter, which allows you to specify the identity, with or without the parameter name.

For identities that have names as well as aliases, you can specify either value as the identity. For example, to retrieve the mailbox object for the user William Stanek with the mail alias Williams, you can use any of the following techniques:

```
get-mailbox Williams
get-mailbox -identity williams
get-mailbox "William Stanek"
get-mailbox -identity 'William Stanek'
```

Typically, Get cmdlets return an object set containing all related items when you omit the identity. For example, if you enter `get-mailbox` without specifying an identity, PowerShell displays a list of all mailboxes available (up to the maximum permitted to return in a single object set).

Cmdlets can display output in several different formats. Although all cmdlets return data in table format by default, there are often many more columns of data than fit across the screen. For this reason, you might need to output data in list format.

To output in list format, redirect the output using the pipe symbol (`|`) to the `Format-List` cmdlet, as shown in this example:

```
get-mailbox "William Stanek" | format-list
```

Because `fl` is an alias for `Format-List`, you also can use `fl`, as in this example:

```
get-mailbox "William Stanek" | fl
```

With a list format output, you should see much more information about the object or the result set than if you were retrieving table-formatted data.

Note also the pipe symbol (`|`) used in the examples. When you are working with Windows PowerShell, you'll often need to use the pipe symbol (`|`) to redirect the output of one cmdlet and pass it as input to another cmdlet. For example, access to remote PowerShell is a privilege for an online user that can be viewed with `Get-User` and managed with `Set-User`. To determine whether a particular user has remote shell access, you can enter:

```
Get-User UserID | fl RemotePowerShellEnabled
```

where `UserID` is the identity of the user to view, such as:

```
Get-User WilliamS | fl RemotePowerShellEnabled
```

If the user should have remote PowerShell access but doesn't currently, you can enable access using the `-RemotePowerShellEnabled` parameter of `Set-User`, as shown in this example:

```
Set-User WilliamS -RemotePowerShellEnabled $true
```

If the user has remote PowerShell access but shouldn't, you can disable access by setting the `-RemotePowerShellEnabled` to `$false`, as shown in this example:

```
Set-User TonyG -RemotePowerShellEnabled $false
```

When you work with list- or table-formatted data, you may want to specify the exact data to display. For example, with `Get-User`, you can display only the user name, display name and remote PowerShell status using:

```
Get-User | Format-Table Name, DisplayName,  
RemotePowerShellEnabled
```

If your organization has a lot of users you can prevent the result set from getting truncated by allowing an unlimited result set to be returned, as shown in this example:

```
Get-User -ResultSize Unlimited | Format-Table  
Name,DisplayName,RemotePowerShellEnabled
```

With cmdlets that have many properties, you may want to filter the output based on a specific property. For example, to display a list of all users who have remote PowerShell access, you can filter the result set on the `RemotePowerShellEnabled` property, as shown in the following example:

```
Get-User -ResultSize unlimited -Filter {RemotePowerShellEnabled -eq $true}
```

Alternatively, you may want to see a list of users who don't have remote PowerShell access. To do this, filter the results by looking for users who have the `RemotePowerShellEnabled` property set to `$False`:

```
Get-User -ResultSize unlimited -Filter {RemotePowerShellEnabled -eq $false}
```

Chapter 3. Getting Started with Users and Contacts

User and contact management is a key part of Exchange administration. User accounts enable individual users to log on to the network and access network resources. In Active Directory, users are represented by User and InetOrgPerson objects.

User objects represent standard user accounts; InetOrgPerson objects represent user accounts imported from non-Microsoft Lightweight Directory Access Protocol (LDAP) or X.500 directory services. User and InetOrgPerson are the only Active Directory objects that can have Exchange mailboxes associated with them.

In contrast, contacts, are people who you or others in your organization want to get in touch with. Contacts can have street addresses, phone numbers, fax numbers, and email addresses associated with them. Unlike user accounts, contacts don't have network logon privileges.

Working with Users and Contacts

In Active Directory, users are represented as objects that can be mailbox-enabled or mail-enabled. A *mailbox-enabled* user account has an Exchange mailbox associated with it. Mailboxes are private storage areas for sending and receiving mail. A user's display name is the name Exchange presents in the global address list.

Another important identifier for mailbox-enabled user accounts is the Exchange alias. The alias is the name that Exchange associates with the account for addressing mail. When your mail client is configured to use Microsoft Exchange Server, you can type the alias or display name in the To, Cc, or Bcc text boxes of an email message and have Exchange Server resolve the alias or name to the actual email address.

Although you'll likely configure most Windows user accounts as mailbox-enabled, user accounts don't have to have mailboxes associated with them. You can create user accounts without assigning a mailbox. You can also create user accounts that are *mail-enabled* rather than mailbox-enabled, which means that the account has an off-site email address associated with it but doesn't have an actual mailbox. Mail-enabled users have Exchange aliases and display names that Exchange Server can resolve to actual email addresses. Internal users can send a message to the mail-enabled user account using the Exchange display name or alias, and the message will be directed to the external address. Users outside the organization can use the Exchange alias to send mail to the user.

It's not always easy to decide when to create a mailbox for a user. To better understand the decision-making process, consider the following scenario:

1. You've been notified that two new users, Elizabeth and Joe, will need access to the domain.
2. Elizabeth is a full-time employee who starts on Tuesday. She'll work on site and needs to be able to send and receive mail. People in the company need to be able to send mail directly to her.
3. Joe, on the other hand, is a consultant who is coming in to help out temporarily. His agency maintains his mailbox, and he doesn't want to have to check mail in two places. However, people in the company need to be able to contact him, and he wants to be sure that his external address is available.
4. You create a mailbox-enabled user account for Elizabeth. Afterward, you create a mail-enabled user account for Joe, ensuring that his Exchange information refers to his external email address.

Mail-enabled users are one of several types of custom recipients that you can create in Exchange Server. Another type of custom recipient is a *mail-enabled* contact. You create a mail-enabled contact so that users can more easily send email to that contact. A mail-enabled contact has an external email address.

Microsoft Exchange Server 2016 has in-place archiving for user mailboxes, which is

designed to replace the need for personal stores in Outlook. An in-place archive is an alternative storage location for historical message data that is seamlessly accessible to a user in Microsoft Outlook 2010 or later and Outlook Web App.

The in-place archive is created as an additional mailbox and is referred to as an archive mailbox. Users can easily move and copy mail data between a primary mailbox and an archive mailbox. Because in-place archiving is a premium feature, an enterprise license is required for each user with an archive mailbox. For more information, see “Working with Archive Mailboxes” in Chapter 6 “Adding Special-Purpose Mailboxes.”

How Email Routing Works: The Essentials

Exchange uses email addresses to route messages to mail servers inside and outside the organization. When routing messages internally, Mailbox servers use mail connectors to route messages to other Exchange servers, as well as to other types of mail servers that your company might use. Two standard types of connectors are used:

- **Send connectors** Control the flow of outbound messages
- **Receive connectors** Control the flow of inbound messages

Send and Receive connectors use Simple Mail Transfer Protocol (SMTP) as the default transport and provide a direct connection among Mailbox servers in an on-premises Exchange organization. Edge Transport servers can also receive mail from and send mail to other types of mail servers.

You can use these connectors to connect Mailbox servers in an organization. When routing messages outside the company, Mailbox servers and Edge Transport servers use mail gateways to transfer messages. The default gateway is SMTP.

Online-only deployments work in much the same way, except that mail is routed through the Exchange Online organization. Here, Exchange Online Protection handles transport.

In hybrid deployments, mailboxes can reside in the on-premises Exchange organization and in an Exchange Online organization. Messages are sent between the organizations transparently and appear as internal messages. To enhance security, messages are encrypted and transferred between the organizations using Transport Layer Security (TLS).

Exchange Server 2016 uses directory-based recipient resolution for all messages that are sent from and received by users throughout an Exchange organization. The Exchange component responsible for recipient resolution is the Categorizer. The Categorizer must be able to associate every recipient in every message with a corresponding recipient object in Active Directory.

All senders and recipients must have a primary SMTP address. If the Categorizer discovers a recipient that does not have a primary SMTP address, it will determine what the primary SMTP address should be or replace the non-SMTP address. Replacing a non-SMTP address involves encapsulating the address in a primary SMTP address that will be used while transporting the message.

IMPORTANT Non-SMTP email address formats include fax, X.400, and the legacy Exchange format (EX). The Categorizer encapsulates email addresses using non-SMTP formats in the Internet Mail Connector Encapsulated Addressing (IMCEA) format. For example, the Categorizer encapsulates the fax address, FAX:888-555-1212, as IMCEA-FAX-888-555-1212@yourdomain.com. Any email address that is longer than what SMTP allows is transmitted as an extended property in the XExch50 field, provided the name part of the address and domain

part of the address don't exceed the allowed limits. The maximum allowed length for an email address in Exchange is 571 characters, 315 characters for the name part of the address, 255 characters for the domain name, and the @ sign character that separates the two name parts.

In addition to primary SMTP email addresses, you can configure alternative recipients and forwarding addresses for users and public folders. If there is an alternative recipient or forwarding address, redirection is required during categorization. You specify the addresses to which messages will be redirected in Active Directory, and redirection history is maintained with each message.

Managing Recipients: The Fundamentals

Exchange Management Shell provides many commands for working with mailbox-enabled users, mail-enabled users, and contacts. The main commands you'll use are shown in the following list:

MAILBOX-ENABLED USER	MAIL-ENABLED USERS	CONTACTS
Connect-Mailbox	Disable-MailUser	Disable-MailContact
Disable-Mailbox	Enable-MailUser	Enable-MailContact
Enable-Mailbox	Get-MailUser	Get-MailContact
Get-Mailbox	New-MailUser	New-MailContact
New-Mailbox	Remove-MailUser	Remove-MailContact
Remove-Mailbox	Set-MailUser	Set-MailContact
Set-Mailbox		


Because Exchange organizations can be on-premises, online, or a hybrid of the two, working with recipients is more complex than it used to be, especially when it comes to creating recipients. Normally, to work with the recipient you access the organization where the recipient should be or has been created. For example, if a mailbox was created in the on-premises Exchange organization, you connect to the on-premises organization and work with the mailbox using the on-premises implementation of Exchange Admin Center or Exchange Management Shell. If a mailbox was created in the online Exchange organization, you connect to the online organization and work with the mailbox using the online implementation of Exchange Admin Center or Exchange Management Shell.

With hybrid deployments, however, you can synchronize users from on-premises Active Directory to Exchange Online. You do this using the hybrid deployment tools. When you run the sync tool for the first time, it copies all of the user accounts, contacts, and groups from Active Directory to Exchange Online. The domains in your organization are then synchronized automatically, so you need to re-run the sync tool only if you add, remove, or rename domains.

Although accounts for synced users are created in the Exchange Online organization,

they are not activated for online use, which means they don't have access to the online features and also haven't been licensed. If you want to create an online mailbox for a synced user, you also must activate the account before the grace period expires. If the user has a local mailbox and you want to move it to Exchange Online, you run the Mailbox Migration Wizard. This wizard configures forwarding of the user's local mailbox to Exchange Online and then copies the user's mailbox data to Exchange Online. Moving and migrating mailboxes is discussed in more detail in Chapter 7 “Managing Mailboxes.”

To create a new synced mailbox user, you have several options. One option is as follows:

1. Create the user account in Active Directory Users And Computers.
2. Wait for the account to be synchronized with Exchange Online.
3. Access the Exchange Online organization. Next, either create the mailbox for the user or migrate the user's existing mailbox to Exchange Online. If you create a mailbox for the user, keep the following in mind:
 - For Exchange Admin Center, this means using the online console for administration. In a synchronized hybrid deployment, you can access the online console from an on-premises console. Click the **Office 365** option. After your browser connects to Office.com, click **Sign In** on the Navigation bar and then select **Work, School Or University** as your account type. Next, provide the email address and password for your Microsoft account and then click **Sign In** to open Office Admin Center. In Office Admin Center, click **Home** () to display the Navigation menu and then click **Exchange** under the Admin Centers heading. This opens the Exchange Admin Center dashboard.
 - For Exchange Online, you access the Exchange Online organization by establishing a remote session as discussed in "Connecting to Exchange Online Using PowerShell" in Chapter 2 "Working with Exchange Online."
4. Using Office 365 Admin Center, activate the synced user and assign a license. When you assign a license, a mailbox is created automatically.

The second option for creating a new synced mailbox user is to use the New-RemoteMailbox cmdlet. In this method, you access the on-premises Exchange organization in Exchange Management Shell and then use New-RemoteMailbox to create an enabled and synced mailbox user, which means:

- A mail-enabled user is created in on-premises Active Directory.
- An associated mailbox is created in Exchange Online.

NOTE Don't forget, you'll also need to assign the user a mailbox plan.

The basic syntax for the RemoteMailbox cmdlets are as follows:

- **New-RemoteMailbox** Creates a mail-enabled user in on-premises Active Directory and a mailbox in Exchange Online.

```
New-RemoteMailbox -Name CommonName [-Alias ExchangeAlias ]
[-ArbitrationMailbox ModeratorMailbox ] [-Archive <$true
| $false>] [-DisplayName Name ] [-DomainController FullyQualifiedName ]
[-FirstName FirstName ] [-Initials Initials ] [-LastName LastName ]
[-ModeratedBy Moderators ] [-ModerationEnabled <$true | $false>]
[-OnPremisesOrganizationalUnit OUName ] [-OverrideRecipientQuotas <$true |
$false>] [-Password Password ] [-PrimarySmtpAddress SmtpAddress ]
[-RemotePowerShellEnabled <$true |$false>] [-RemoteRoutingAddress
ProxyAddress ] [-ResetPasswordOnNextLogon <$true | $false>]
[-SamAccountName PreWin2000Name ] [-SendModerationNotifications <Never |
Internal | Always>] [-UserPrincipalName LoginName ]
```

- **Enable-RemoteMailbox** Creates an online mailbox for a user already created in on-premises Active Directory.

```
Enable-RemoteMailbox-Identity UserId [-Alias ExchangeAlias ]
[-DisplayName DisplayName ] [-DomainController DomainControllerName ]
[-PrimarySmtpAddress SmtpAddress ] [-RemoteRoutingAddress ProxyAddress ]
```

- **Disable-RemoteMailbox** Removes an online mailbox but keeps the user account in on-premises Active Directory.

```
Disable-RemoteMailbox -Identity UserId [-Archive <$true | $false>]
[-DomainController DomainControllerName ] [-IgnoreDefaultScope<$true |
$false>] [-IgnoreLegalHold <$true | $false>]
```

- **Remove-RemoteMailbox** Removes an online mailbox and the related account in on-premises Active Directory.

```
Remove-RemoteMailbox -Identity UserId [-Archive <$true | $false>]
[-DomainController DomainControllerName ] [-IgnoreDefaultScope<$true |
$false>] [-IgnoreLegalHold <$true | $false>]
```

Regardless of which approach you use to create new mailbox users in Exchange Online, you must license these mailbox users in Office 365. You do this by associating a mailbox plan with each mailbox user. Using the graphical tools, you can associate mailbox plans when you are creating mailbox users or afterward by editing the account properties. In a remote session with Exchange Online, you can use the `-MailboxPlan` parameter with the `New-Mailbox` cmdlet to do the same. However, at the time of this writing, there are no mailbox plan parameters for any of the `RemoteMailbox` cmdlets. (Hopefully, this oversight will be corrected by the time you read this.)

When you assign mailbox plans, you need to ensure you have enough licenses. You purchase and assign licenses using the billing and subscription options in Office 365 Admin Center. Select **Billing** on the dashboard or click **Subscriptions** under the Billing heading in the Navigation menu to see the subscription and licensing options.

Office 365 will allow you to assign more mailbox plans than you have licenses for. However, after the initial grace period, problems will occur. For example, mail data for unlicensed mailboxes may become unavailable. Remember, the number of valid licenses shouldn't exceed the number of assigned licenses.

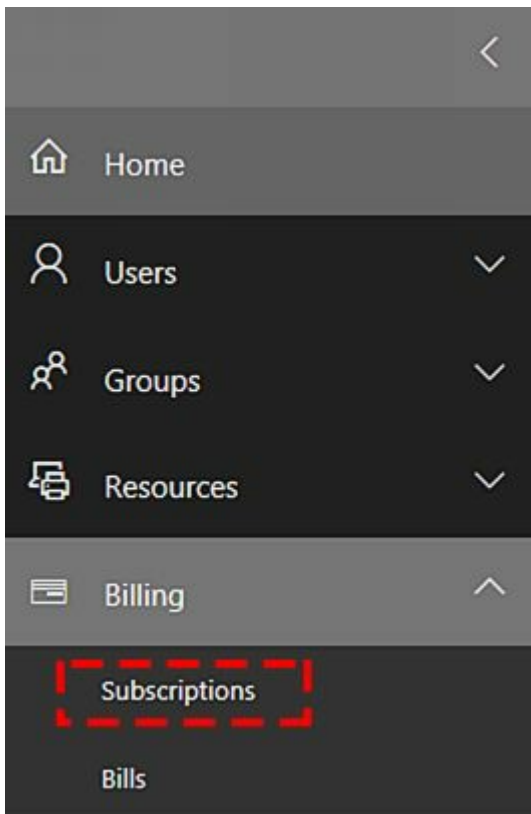


FIGURE 3-1 Accessing subscriptions in Office 365 Admin Center.

You activate and license synced users in Office 365 as well:

1. Select **Users** on the dashboard or click **Active Users** under the Users heading in the Navigation menu to display active users.




FIGURE 3-2 Accessing active users in Office 365 Admin Center.

2. On the Users page, click a user that you want to activate and license to display account settings.
3. Click **Edit** on the Product Licenses panel.
4. As shown in Figure 3-3, specify the work location for the user.
5. Select the mailbox plan to assign by clicking it to the On position.
6. Optionally, click to switch individual plan options on or off.

7. Select **Activate** or **Save** as appropriate.



FIGURE 3-3 Switch the plan on and configure its options.

The Office 365 service, its settings and accounts are all manageable from Windows PowerShell. Every account you create in the online environment is in fact created in the online framework within which Office 365 and Exchange Online operate. This framework is called Windows Azure, and like Windows Server, it uses Active Directory to provide its directory services. You can manage Office 365 from Windows PowerShell as discussed in “Connecting to Windows Azure” in Chapter 2.

Finding Existing Mailboxes, Contacts, And Groups

You work with recipients where they were created, which can be either in an on-premises Exchange organization or in Exchange Online. You can view current mailboxes, mail-enabled users, contacts, and groups by following these steps:

1. Open Exchange Admin Center using one of the following techniques:
 - For on-premises Exchange, open your Web browser and then enter the secure URL for Exchange Admin Center, such as *https://mailserver16.imagedlands.com/ecp* .
 - For online Exchange, open your Web browser and then enter the secure URL for Office 365 Admin Center, such as *https://portal.microsoftonline.com/admin/default.aspx* . In Office 365 Admin Center, click **Exchange** under the Admin Centers heading on the Navigation menu to open the Exchange Online version of Exchange Admin Center.
2. As shown in Figure 3-4, select **Recipients** in the Navigation menu and then select the related Mailboxes, Groups, or Contacts tab, as appropriate for the type of recipient you want to work with.

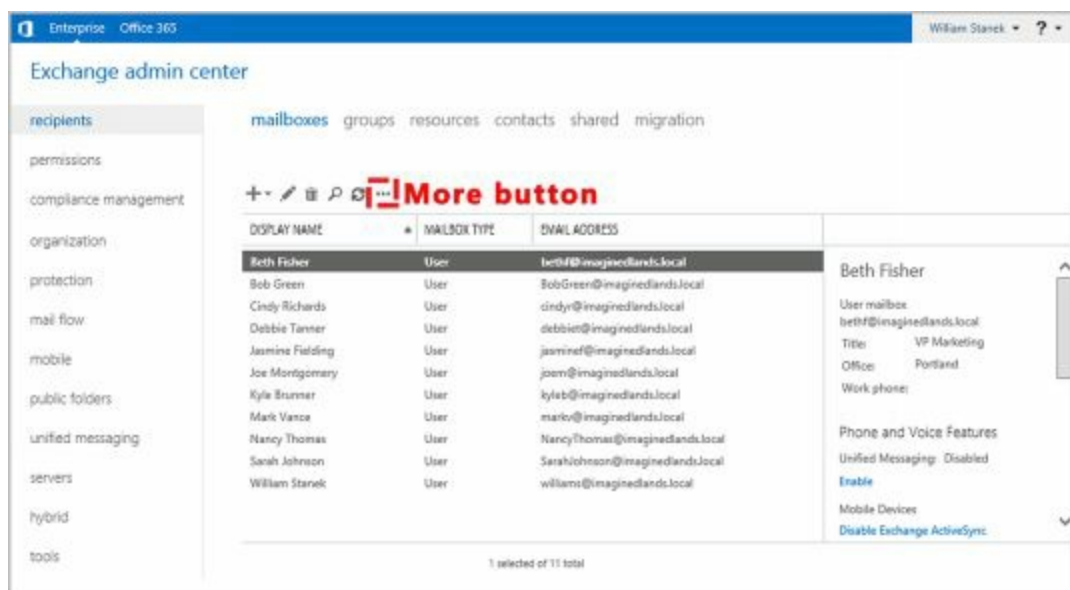


FIGURE 3-4 Accessing the Recipient node to work with mailboxes, distribution groups, and mail contacts.

By default, all recipients of the selected type are displayed. With mailboxes this means that user mailboxes, linked user mailboxes, legacy user mailboxes, and remote user mailboxes are displayed.

By default, Exchange Admin Center displays only three columns of information for each recipient, including the display name, mailbox type, and email address. To customize the columns of information displayed, click the More button (**⋮**) and then select **Add/Remove Columns** . Use the options provided in the Add/Remove Columns dialog box, shown in Figure 3-5, to configure the columns to use, and then click **OK** .

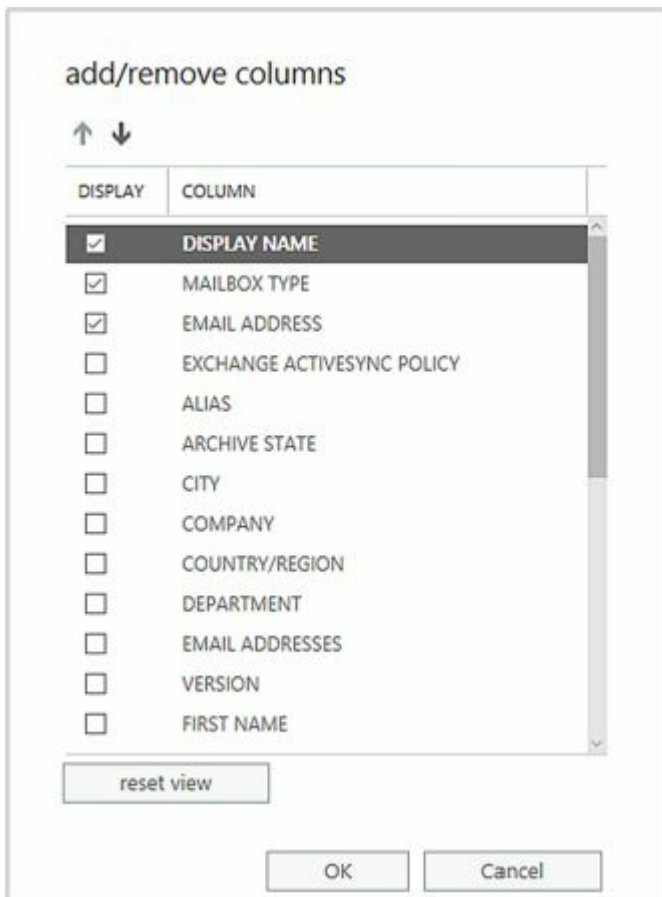


FIGURE 3-5 Customizing the list of columns to display using the options provided.

In large organizations, you may need to filter based on attributes to locate recipients you want to work with. To do this, click the More button (**⋮**) and then select **Advanced Search** . Next, use the Advanced Search dialog box, shown in Figure 3-6 to filter by alias, display name, department, email addresses, first name, last name, and recipient type. The Recipient Types condition allows you to filter the results for specific recipient subtypes, such as only remote mailbox users.

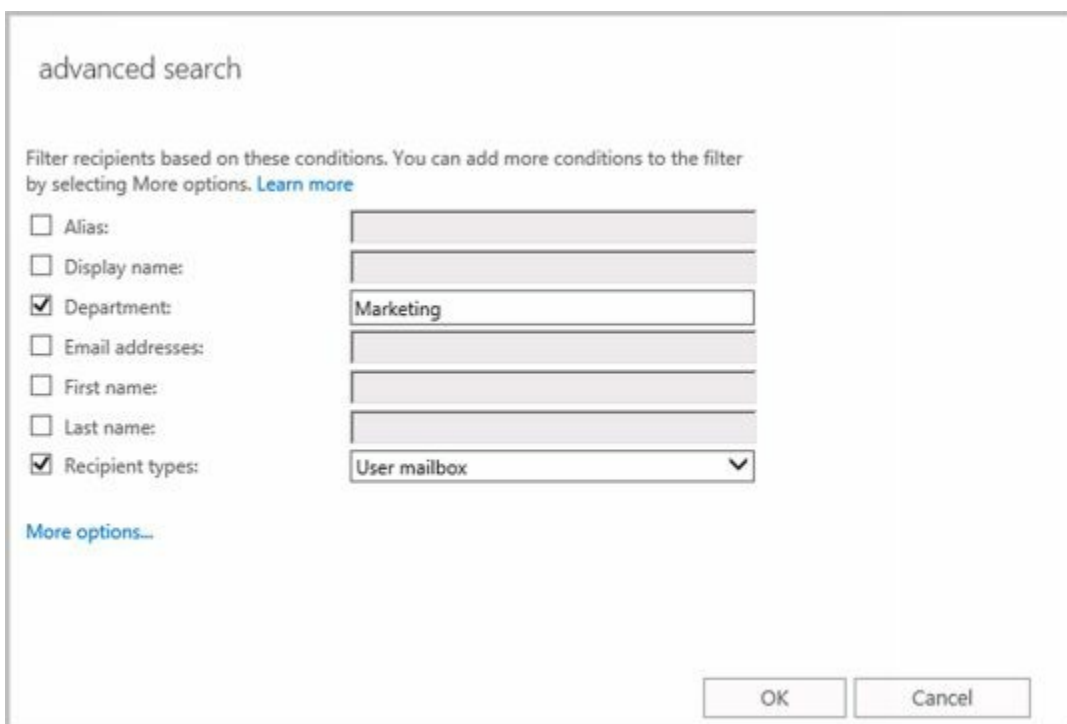


FIGURE 3-6 Performing advanced searches with filters.

You can add conditions that allow you to filter results based on city, state, country, office, title, group membership, and more:

1. Select **More Options** and then select **Add Condition**.
2. Click in the selection list and then select the condition, such as **City**.
3. Type the search word or phrase and then click **OK**.
4. Repeat this process to add other conditions.

In Exchange Management Shell, you can find mailboxes, contacts, and groups by using the following commands:

- **Get-User** Use the **Get-User** cmdlet to retrieve all users in the forest that match the specified conditions.

```
Get-User [-Identity UserId | -Anr Identifier ] [-AccountPartition PartitionId ]  
[-Arbitration <$true | $false>][-Credential Credential ]  
[-DomainController DomainControllerName ] [-Filter FilterString]  
[-IgnoreDefaultScope <$true | $false>] [-Organization OrgName ]  
[-OrganizationalUnit OUName ] [-PublicFolder <$true | $false>]  
[-ReadFromDomainController <$true | $false>] [-RecipientTypeDetails  
Details ] [-ResultSize Size ] [-SortBy String ]
```

- **Get-Contact** Use the **Get-Contact** cmdlet to retrieve information about a specified contact or contacts.

```
Get-Contact [-Identity ContactId | -Anr ContactID ] [-AccountPartition  
PartitionId ][-Credential Credential ] [-DomainController  
DomainControllerName ] [-Filter FilterString ] [-IgnoreDefaultScope <$true  
| $false>] [-Organization OrgName ] [-OrganizationalUnit OUName ]  
[-ReadFromDomainController <$true | $false>]  
[-RecipientTypeDetails Details ] [-ResultSize Size ] [-SortBy Value ]
```

- **Get-Group** Use the **Get-Group** cmdlet to query for existing groups.

```
Get-Group [-Identity GroupId | -Anr GroupID ]  
[-AccountPartition PartitionId ] [-Credential Credential ]  
[-DomainController FullyQualifiedName ] [-Filter FilterString ]  
[-IgnoreDefaultScope <$true | $false>] [-Organization OrgName ]  
[-OrganizationalUnit OUName ] [-ReadFromDomainController <$true |  
$false>] [-RecipientTypeDetails {"Contact" | "MailContact" |  
"MailUser" | "RoleGroup" | "User" | "UserMailbox" | ... }]  
[-ResultSize Size ] [-SortBy Value ]
```

- **Get-RemoteMailbox** Use the **Get-RemoteMailbox** cmdlet to get details for mail-enabled users in on-premises Active Directory that have mailboxes in Exchange Online.

```
Get-RemoteMailbox [-Identity UserId | -Anr Identifier ] [-Alias
```

ExchangeAlias] [-Archive <\$true | \$false>] [-DomainController
DomainControllerName] [-OnPremisesOrganizationalUnit **OUName**]
[-ReadFromDomainController **DomainControllerName**]
[-ResultSize **NumResults**]

Finding Synced, Unlicensed, Inactive, and Blocked Users

When you are working with hybrid organizations, users can be synced from Active Directory to Exchange Online. These synced users can have mailboxes on-premises or in Exchange Online. If you need to view all the synced users, determine where a synced user's mailbox is located, or perform other tasks with synced users, complete the following steps:

1. Open Office 365 Admin Center. Select **Users** in the Navigation menu, and then click **Active Users**.
2. On the Filters drop-down list, select **Synced Users**.
3. You should now see a list of synced users.

A synced user is only one type of user you may want to find in an Exchange Online organization. You also may want to find:

- **Unlicensed users** These users haven't been assigned an Exchange Online license. Although there is a grace period for licensing after creating a mailbox user online, the user may lose mailbox data after the grace period expires.
- **Inactive users** These users have been deleted by an admin, which puts them in inactive status for a period of 30 days. When the recovery period expires, the account and any unprotected data is removed.
- **Sign-in Allowed users** These users can sign in and the related accounts are active.
- **Sign-in Blocked users** These users cannot sign in and the related accounts are blocked, such as may happen when a user's password expires.
- **Users with errors** These users have errors associated with their accounts.

You can find allowed users, blocked users, unlicensed users, or users with errors by completing the following steps:

1. Open Office 365 Admin Center. Select **Users** in the Navigation menu, and then click **Active Users**.
2. On the Filters drop-down list, select Sign-in Allowed Users, Sign-in Blocked Users, Unlicensed Users, or Users With Errors as appropriate.

In Office 365 Admin Center, you can find inactive users by selecting Users in the Navigation menu and then selecting the Deleted Users tab.

Chapter 4. Managing Users

In Exchange Server 2016, Exchange Admin Center and Exchange Management Shell are the primary administration tools you use to manage mailboxes and mail contacts. You can use these tools to create and manage mail-enabled user accounts, mailbox-enabled user accounts, and mail-enabled contacts, as well as any other configurable aspect of Exchange Server.

The sections that follow examine techniques to manage user accounts and the related Exchange features of those accounts whether you are working with either on-premises Exchange organizations or Exchange Online. In a hybrid environment, you always manage domain user accounts and their mailboxes using the on-premises Exchange tools. Your changes are then synced to the online environment.

NOTE Domain administrators can create user accounts and contacts using Active Directory Users And Computers. If any existing user accounts need to be mail-enabled or mailbox-enabled, you perform these tasks using the Exchange management tools. If existing contacts need to be mail-enabled, you also perform this task using the Exchange management tools.

Creating Mailbox-Enabled and Mail-Enabled User Accounts

Generally speaking, you need to create a user account for each user who wants to use network resources. The following sections explain how to create domain user accounts that are either mailbox-enabled or mail-enabled, and how to add a mailbox to an existing user account. If a user needs to send and receive email, you need to create a new mailbox-enabled account for the user or add a mailbox to the user's existing account. Otherwise, you can create a mail-enabled account.

Working with Logon Names and Passwords

Before you create a domain user account, you should think for a moment about the new account's logon name and password. You identify all domain user accounts with a logon name. This logon name can be (but doesn't have to be) the same as the user's email address. In Windows domains, logon names have two parts:

- **User name** The account's text label
- **User domain** The domain where the user account exists

For the user Williams whose account is created in `imaginedlands.com`, the full logon name for Windows is `williams@imaginedlands.com`.

User accounts can also have passwords and public certificates associated with them. *Passwords* are authentication strings for an account. *Public certificates* combine a public and private key to identify a user. You log on with a password by typing the password. You log on with a public certificate by using a smart card and a smart card reader.

Although Windows displays user names to describe privileges and permissions, the key identifiers for accounts are security identifiers (SIDs). SIDs are unique identifiers that Windows generates when you create accounts. SIDs consist of the domain's security ID prefix and a unique relative ID. Windows uses these identifiers to track accounts independently from user names. SIDs serve many purposes; the two most important are to allow you to easily change user names and to allow you to delete accounts without worrying that someone could gain access to resources simply by re-creating an account with the same user name.

When you change a user name, you tell Windows to map a particular SID to a new name. When you delete an account, you tell Windows that a particular SID is no longer valid. Afterward, even if you create an account with the same user name, the new account won't have the same privileges and permissions as the previous one because the new account will have a new SID.


Mail-Enabling New User Accounts

Mail-enabled users are defined as custom recipients in Exchange Server. They have an Exchange alias and an external email address, but they do not have an Exchange mailbox. All email messages sent to a mail-enabled user are forwarded to the remote email address associated with the account.

In Exchange Admin Center, mail-enabled users are listed as Mail Users under Recipients > Contacts. You can manage mail-enabled users through Exchange Admin Center and Exchange Management Shell.

NOTE With on-premises Exchange, you have two options for mail-enabled users and contacts that are no longer needed. You can disable the mail-enabled user or contact, or you can delete the mail-enabled user or contact. With Exchange online, your only option is to delete the mail-enabled user or contact.

You can create a new mail-enabled user by completing the following steps:

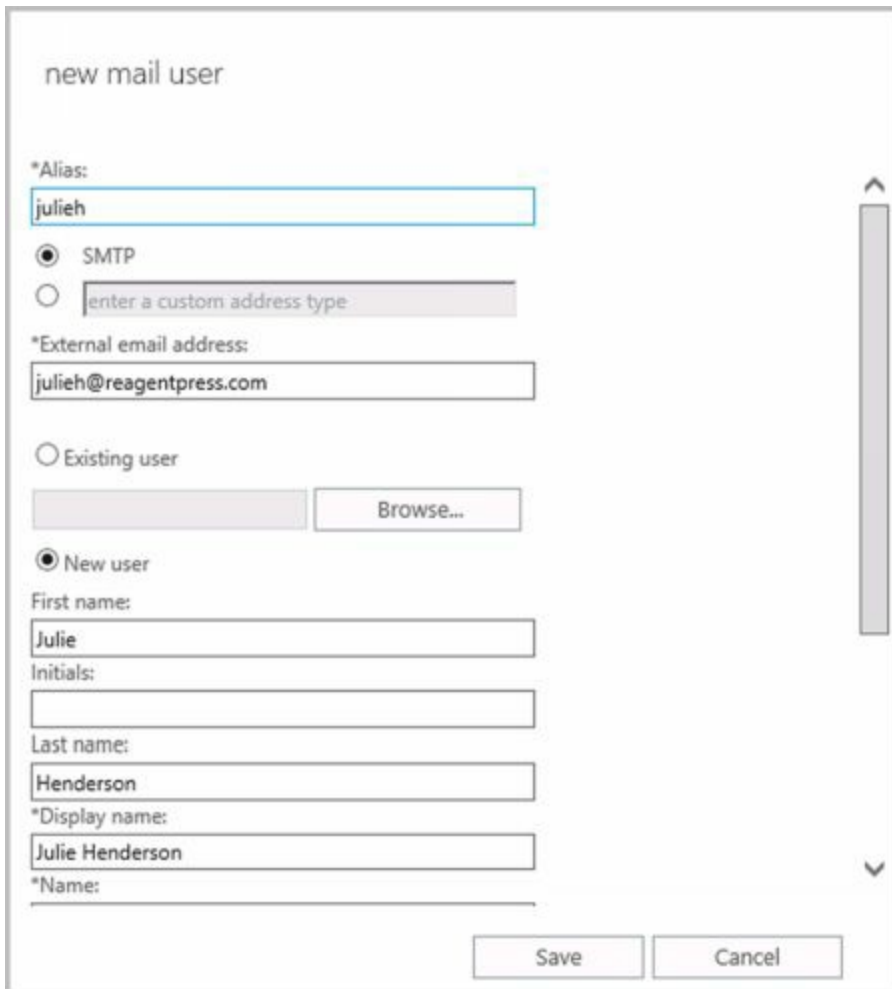
1. In Exchange Admin Center, select **Recipients** in the Navigation menu and then select **Contacts**.
2. Click **New** () and then select **Mail User**. This opens the New Mail User dialog box, shown in Figure 4-1.
3. If you are working with on-premises Exchange, Existing User is selected by default and you'll need to select **New User** instead.
4. Type the user's first name, middle initial, and last name in the text boxes provided. These values are used to create the Display Name entry (as well as the Active Directory name with on-premises Exchange).
5. The Display Name and Name properties can't exceed 64 characters. As necessary, make changes to the Display Name, Name, or both text boxes. For example, you might want to type the name in LastName FirstName MiddleInitial format or in FirstName MiddleInitial LastName format.

IMPORTANT The difference between the Display Name and the Name properties is subtle but important. The Display Name property sets the name displayed in Exchange and Outlook. The Name property sets the display name in Active Directory and is the Common Name (CN) value associated with the user.

6. In the Alias text box, type an alias for the mail-enabled user. This alias should uniquely identify the mail-enabled user in the Exchange organization. Alias names cannot contain spaces.
7. In the External Email Address text box, type the mail user's external email address. By default, the address is configured as a standard SMTP email address. If you are working with on-premises Exchange, you can specify a custom address type by selecting the related option and then entering a prefix that identifies the custom type. Use X.400, GroupWise or Lotus Notes for X.400, GroupWise and Lotus Notes address types respectively.
8. With on-premises Exchange, the user account is created in the default user

container, which typically is the Users container. Because you'll usually need to create new user accounts in a specific organizational unit rather than in the Users container, click **Browse** to the right of the Organizational Unit text box. In the Select Organizational Unit dialog box, choose the location where you want to store the account and then click **OK**.

9. In the User ID or User Logon Name text box, type the user's logon name. Use the drop-down list to select the domain with which you want to associate the account. This sets the fully qualified logon name, such as `williams@imaginedlands.com`.



The screenshot shows a 'new mail user' configuration dialog box. The title is 'new mail user'. The fields and options are as follows:

- *Alias: julieh
- SMTP
- enter a custom address type
- *External email address: julieh@reagentpress.com
- Existing user (with a 'Browse...' button)
- New user
- First name: Julie
- Initials: (empty)
- Last name: Henderson
- *Display name: Julie Henderson
- *Name: (empty)

At the bottom, there are 'Save' and 'Cancel' buttons.

FIGURE 4-1 Configuring the mail-enabled user's settings.

10. Type and then confirm the password for the account. This password must follow the conventions of your organization's password policy. Typically, this means that the password must include at least eight characters and must use three of the four available character types: lowercase letters, uppercase letters, numbers, and symbols.
11. With on-premises Exchange you can select Require Password Change On Next Logon check box to ensure that the user changes the password at next logon.
12. Click **Save**. Exchange Admin Center creates the new mail-enabled user.

If an error occurs, the user will not be created. You will need to click **OK**, correct the problem, and then click **Save** again. Consider the error example shown in Figure 4-2. In

this instance, the user logon name/user ID was already in use so the user couldn't be created.

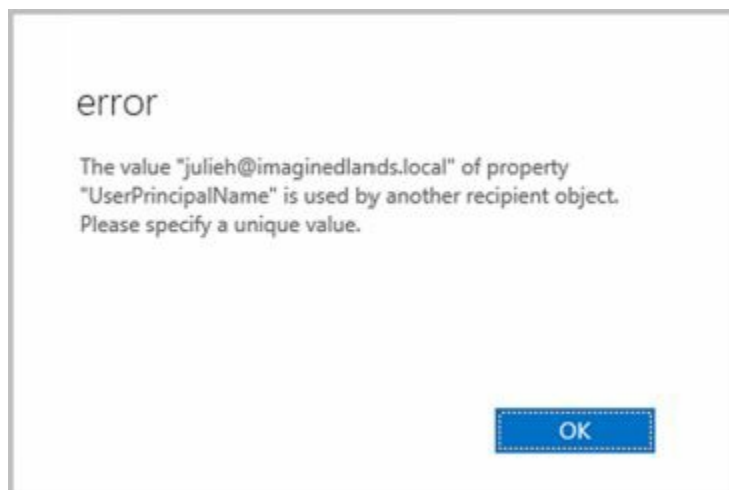


FIGURE 4-2 An error occurs when a user's principal name is already in use.

You can list all mail-enabled users by typing **get-mailuser** at the Exchange Management Shell prompt. Sample 4-1 provides the full syntax and usage for Get-MailUser.

SAMPLE 4-1 Get-MailUser cmdlet syntax and usage

Syntax

```
Get-MailUser [-Identity Identifier | -Anr Name ] [-AccountPartition PartitionId ]  
[-Credential Credential ] [-DomainController FullyQualifiedName ] [-Filter  
FilterString ] [-IgnoreDefaultScope {$true | $false}] [-Organization OrgName ]  
[-OrganizationalUnit OUName ] [-ReadFromDomainController {$true | $false}]  
[-ResultSize Size ] [-SortBy Value ]
```

Usage

```
Get-MailUser -Identity "aaronl" | fl
```

```
Get-MailUser -OrganizationalUnit "marketing" | fl
```

NOTE By default, Get-MailUser lists the name and recipient type for matches. In the example, fl is an alias for Format-List and is used to get detailed information about matching entries.

You can create a new mail-enabled user account using the New-MailUser cmdlet. Sample 4-2 shows the syntax and usage. When prompted, provide a secure password for the user account.

NOTE The syntax and usage are entered on multiple lines for ease of reference. You must enter the command-line values for a cmdlet on a single line.

SAMPLE 4-2 New-MailUser cmdlet syntax and usage

Syntax

```
New-MailUser -Name CommonName -ExternalEmailAddress EmailAddress  
[-Password Password ] [-UserPrincipalName LoginName ] {AddtlParams1}
```

New-MailUser -Name **CommonName** -FederatedIdentity **FederatedId**
-WindowsLiveID **WindowsLiveId** [-EvictLiveId <\$true | \$false>]
[-ExternalEmailAddress **EmailAddress**] [-NetID **NetID**] {AddtlParams2}

New-MailUser -Name **CommonName** -FederatedIdentity **FederatedId**
-MicrosoftOnlineServicesID **WindowsLiveId** [-NetID **NetID**] {AddtlParams2}

New-MailUser -Name **CommonName** -ImportLiveId <\$true | \$false>
-WindowsLiveID **WindowsLiveId** [-ExternalEmailAddress **EmailAddress**]
[-UsageLocation **CountryInfo**] {AddtlParams2}

New-MailUser -Name **CommonName** [-MicrosoftOnlineServicesID **WindowsLiveId**]
{AddtlParams2}

New-MailUser -Name **CommonName** -MicrosoftOnlineServicesID **WindowsLiveId**
-Password **Password** [-ExternalEmailAddress **EmailAddress**] [-UsageLocation
CountryInfo] {AddtlParams2}

New-MailUser -Name **CommonName** -Password **Password** -WindowsLiveID
WindowsLiveId [-EvictLiveId <\$true | \$false>] [-ExternalEmailAddress
EmailAddress] [-UsageLocation **CountryInfo**] {AddtlParams2}

New-MailUser -Name **CommonName** -UseExistingLiveId <\$true | \$false>
-WindowsLiveID **WindowsLiveId** [-BypassLiveId <\$true | \$false>]
[-ExternalEmailAddress **EmailAddress**] [-NetID **NetID**]
[-UsageLocation **CountryInfo**] {AddtlParams2}

{AddtlParams1}
[-Alias **ExchangeAlias**] [-ArbitrationMailbox **ModeratorMailbox**]
[-DisplayName **Name**] [-DomainController **FullyQualifiedName**] [-FirstName
FirstName] [-Initials **Initials**] [-LastName **LastName**]
[-MacAttachmentFormat <BinHex | UuEncode | AppleSingle | AppleDouble>]
[-MessageBodyFormat <Text | Html | TextAndHtml>] [-MessageFormat <Text |
Mime>] [-ModeratedBy **Moderators**] [-ModerationEnabled <\$true | \$false>]
[-Organization **OrgName**] [-OrganizationalUnit **OUName**] [-PrimarySmtpAddress
} **SmtpAddress**] [-ResetPasswordOnNextLogon <\$true | \$false>]
} [-SamAccountName **PreWin2000Name**] [-SendModerationNotifications <Never |
Internal | Always>] [-UsageLocation **CountryInfo**] [-UsePreferMessageFormat
<\$true | \$false>]

{AddtlParams2}
[-Alias **ExchangeAlias**] [-ArbitrationMailbox **ModeratorMailbox**]
[-DisplayName **Name**] [-DomainController **FullyQualifiedName**] [-FirstName
FirstName] [-Initials **Initials**] [-LastName **LastName**] [-ModeratedBy
Moderators] [-ModerationEnabled <\$true | \$false>] [-Organization **OrgName**]
[-OrganizationalUnit **OUName**] [-PrimarySmtpAddress **SmtpAddress**]
[-RemotePowerShellEnabled <\$true | \$false>] [-ResetPasswordOnNextLogon

<\$true | \$false>] [-SamAccountName **PreWin2000Name**]
 [-SendModerationNotifications <Never | Internal | Always>]

Usage

```
New-MailUser -Name "Frank Miller" -Alias "Frankm"  

  -OrganizationalUnit "imaginedlands.local/Technology"  

  -UserPrincipalName "Frankm@imaginedlands.local" -SamAccountName "Frankm"  

  -FirstName "Frank" -Initials "" -LastName "Miller"  

  -ResetPasswordOnNextLogon $false  

  -ExternalEmailAddress "SMTP:Frankm@hotmail.com"
```

Mail-Enabling Existing User Accounts

When a user already has an account in Active Directory, you can mail-enable the account using Exchange Admin Center and Exchange Management Shell. In Exchange Admin Center for your on-premises organization, you can mail-enable an existing user account by completing the following steps:


1. Select **Recipients** in the Navigation menu and then select **Contacts** .
2. Click **New** () and then select **Mail User** . This opens the New Mail User dialog box.
3. In the Alias text box, type an alias for the mail-enabled user. This alias should uniquely identify the mail-enabled user in the Exchange organization. Alias names cannot contain spaces.
4. In the External Email Address text box, type the mail user's external email address. By default, the address is configured as a standard SMTP email address. If you are working with on-premises Exchange, you can specify a custom address type by selecting the related option and then entering a prefix that identifies the custom type. Use X.400, GroupWise or Lotus Notes for X.400, GroupWise and Lotus Notes address types respectively.
5. The Existing User option is selected by default, as shown in Figure 4-3. Click **Browse** . This displays the Select User dialog box.

FIGURE 4-3 Configuring mail for an existing user.

6. In the Select User dialog box, select the user account you want to mail-enable and then click **OK**. User accounts that are not yet mail-enabled or mailbox-enabled for the current domain are listed by name and organizational unit.
7. Click **Save**. Exchange Admin Center mail-enables the user account you previously selected. If you're working in a synced, hybrid organization, the mail-enabled user will be synced to Exchange Online as well. If an error occurs, the user account will not be mail-enabled. You will need to correct the problem and repeat this procedure. Click **Finish**.

You can mail-enable an existing user account using the Enable-MailUser cmdlet. Sample 4-3 shows the syntax and usage. For the identity parameter, you can use the user's display name, logon name, or user principal name.

SAMPLE 4-3 Enable-MailUser cmdlet syntax and usage

Syntax

```
Enable-MailUser -Identity Identity -ExternalEmailAddress EmailAddress
[-Alias ExchangeAlias ] [-DisplayName Name ] [-DomainController
FullyQualifiedName ] [-MacAttachmentFormat <BinHex | UuEncode |
AppleSingle | AppleDouble>] [-MessageBodyFormat <Text | Html |
TextAndHtml>] [-MessageFormat <Text | Mime>] [-PrimarySmtpAddress
```

SmtpAddress] [-UsePreferMessageFormat <\$true | \$false>]

Usage

```
Enable-MailUser -Identity "imaginedlands.local/Marketing/Frank Miller"  
-Alias "Frankm" -ExternalEmailAddress "SMTP:Frankm@hotmail.com"
```

Managing Mail-Enabled User Accounts

You can manage mail-enabled users in several ways. If a user account should no longer be mail-enabled, you can disable mail forwarding. To disable mail forwarding in Exchange Admin Center for your on-premises organization, select Recipients in the Navigation menu and then select the Contacts tab. Next, select the user you want to disable. Click the More button (**...**) and then select **Disable**. When prompted to confirm, select **Yes**. If you're working in a synced, hybrid organization, this change will be synced to Exchange Online as well.

At the Exchange Management Shell prompt, you can disable mail forwarding using the Disable-MailUser cmdlet, as shown in Sample 4-4.

SAMPLE 4-4 Disable-MailUser cmdlet syntax and usage

Syntax

```
Disable-MailUser -Identity Identity [-DomainController  
FullyQualifiedName ] [-IgnoreDefaultScope {$true | $false}]
```

Usage

```
Disable-MailUser -Identity "Frank Miller"
```

If you no longer need a mail-enabled user account, you can permanently remove it from Active Directory. To remove a mail-enabled user account in Exchange Admin Center for your on-premises organization, select the mail user and then select the Delete option. When prompted to confirm, click Yes. If you're working in a synced, hybrid organization, this change will be synced to Exchange Online as well.

At the Exchange Management Shell prompt, you can remove a mail-enabled user account by using the Remove-MailUser cmdlet, as shown in Sample 4-5.

SAMPLE 4-5 Remove-MailUser cmdlet syntax and usage

Syntax


```
Remove-MailUser -Identity "Identity" [-DomainController DCName ]  
[-IgnoreDefaultScope {$true | $false}]  
[ -KeepWindowsLiveID {$true | $false}]
```

Usage

```
Remove-MailUser -Identity "Frank Miller"
```

Creating Domain User Accounts with Mailboxes

You can create a new domain user account with a mailbox in several ways. If you are using a hybrid configuration and want the user created in Active Directory and the mailbox created in Exchange online, you can use the techniques discussed earlier under "Understanding on-premises and online recipient management." Otherwise, you can create a new domain user account and a mailbox for that account using only your on-premises Exchange administration tools. To do this, complete the following steps:

1. In Exchange Admin Center, select **Recipients** in the Navigation menu and then select **Mailboxes**.
2. Click **New** () and then select **User Mailbox**. This opens the New User Mailbox dialog box, shown in Figure 4-4.
3. In the **Alias** text box, type an alias for the mailbox user. This alias should uniquely identify the user in the Exchange organization. Alias names cannot contain spaces.

NOTE The alias and domain suffix are combined to create the email address for the user. For example, if the alias is tedc and the domain suffix is imagedlands.com, the email address is set as tedc@imagedlands.com.

4. Select **New User**. Type the user's first name, middle initial, and last name in the text boxes provided. These values are used to create the Display Name entry as well as the Active Directory name with on-premises Exchange.

FIGURE 4-4 Configuring the mailbox user's settings.

5. The Display Name and Name properties can't exceed 64 characters. As necessary, make changes to the Display Name, Name, or both text boxes. For example, you might want to type the name in LastName FirstName MiddleInitial format or in FirstName MiddleInitial LastName format.

IMPORTANT The difference between the Display Name and the Name properties is subtle but important. The Display Name property sets the name displayed in Exchange and Outlook. The Name property sets the display name in Active Directory and is the Common Name (CN) value associated with the user.

6. Unless you specify otherwise, the user account is created in the default user container, which typically is the Users container. Because you'll usually need to create new user accounts in a specific organizational unit rather than in the Users container, click **Browse** to the right of the Organizational Unit text box. In the Select An Organizational Unit dialog box, shown in Figure 4-5, choose the location to store the account and then click **OK**.



FIGURE 4-5 Selecting the organizational unit for the new user.

7. In the User Logon Name text box, type the user's logon name. Use the drop-down list to select the domain with which you want to associate the account. This sets the fully qualified logon name, such as jeffp@imaginedlands.local.
8. Type and then confirm the password for the account. This password must follow the conventions of your organization's password policy. Typically, this means that the password must include at least eight characters and must use three of the four available character types: lowercase letters, uppercase letters, numbers, and symbols.
9. You can select the Require Password Change On Next Logon check box to ensure that the user changes the password at next logon.
10. Click **More Options** . At this point, you can do the following:
 - **Specify the mailbox database** Exchange uses the mailbox provisioning load balancer to select a database to use when you create a mailbox and do not specify the mailbox database to use. If you want to specify the database to use, click Browse to the right of the Mailbox Database box. In the Select Mailbox Database dialog box, you'll see a list of available mailbox databases listed by name, server, and Exchange version. Select the mailbox database to use and then select OK.
 - **Create an archive mailbox** If you want to create an archive mailbox for the user,

select the related check box. Items in the user's mailbox will be moved automatically to the archive mailbox based on the default retention policy. Using the related Browse option, you also can choose a mailbox database for the archive. If you don't choose a mailbox database for the archive, Exchange chooses one for you.

- **Assign an address book policy** By default, a user has access to the full address book information in the organization. Using address book policies, you can create customized address books. To apply an available policy, select it from the drop-down list.

11. Click **Save** . Exchange Admin Center creates the new mailbox user. If an error occurs, neither the user nor the mailbox will be created. You will need to click OK, correct the problem, and then click Save again.

Creating the user account and mailbox isn't necessarily the final step. You might also want to do the following:

- Add detailed contact information for the user, such as a business phone number and title
- Add the user to security and distribution groups
- Enable or disable mailbox features for the account
- Modify the user's default delivery options, storage limits, and restrictions on the account
- Associate additional email addresses with the account

NOTE For all mailbox-enabled accounts, an SMTP email address is configured automatically. You can also add more addresses of the same type. For example, if Brian Johnson is the company's human resources administrator, he might have the primary SMTP address of `brianj@imaginedlands.com` and an alternate SMTP address of `resumes@imaginedlands.com`.

You may also want to apply appropriate policies to the mailbox. Various types of policies control how users access their mailboxes and how mailbox data is stored. These policies include:

- **Address book policy** Controls access to the address book information in the organization and allows you to create custom views for various users. A default address book policy is not created when you install Exchange 2016. You can check to see if any address book policies have been created by entering `get-addressbookpolicy` in Exchange Management Shell.
- **Mobile device mailbox policy** Controls security settings for mobile devices. When you install Exchange Server, a default mobile device mailbox policy is created and applied automatically to all new mailboxes you create unless you specify a different policy to use. To view the settings for the default policy, enter `get-mobiledevicemailboxpolicy -identity "Default"` in Exchange Management Shell.
- **Retention policy** Specifies the delete and move-to-archive rules that are applied to items in mailboxes. Exchange Server 2016 uses retention policies and retention tags as part of the Messaging Records Management feature. When you install Exchange

2016a default retention policy is created but is not applied to new mailboxes by default. Therefore, you must explicitly assign a retention policy. To view the settings for the default policy, enter **get-retentionpolicy –identity “Default MRM Policy” | fl** in Exchange Management Shell.

- **Role assignment policy** Controls management roles assigned to users. When you install Exchange Server, a default role assignment policy is created and applied automatically to all new mailboxes you create unless you specify a different policy to use. To view the settings for the default policy, enter **get-roleassignmentpolicy –identity “Default Role Assignment Policy”** in Exchange Management Shell.
- **Sharing policy** Controls how users can share calendar and contact information with users outside your organization. When you install Exchange Server, a default sharing policy is created and applied automatically to all new mailboxes you create unless you specify a different policy to use. To view the settings for the default policy, enter **get-sharingpolicy –identity “Default Sharing Policy”** in Exchange Management Shell.

In Exchange Management Shell, you can create a user account with a mailbox by using the **New-Mailbox** cmdlet. Sample 4-6 provides the syntax and usage. When you are prompted, enter a secure password for the new user account.

SAMPLE 4-6 New-Mailbox cmdlet syntax and usage

Syntax

```
New-Mailbox -Name Name -Password Password -UserPrincipalName  
UserNameAndSuffix {AddtlParams} {CommonParams} {ModParams}
```

```
New-Mailbox -Name Name -Room <$true | $false>[-Office OfficeName ]  
[-Password Password ] [-Phone PhoneNumber ] [-ResourceCapacity Capacity ]  
[-UserPrincipalName UserNameAndSuffix ] {CommonParams} {ModParams}
```

```
New-Mailbox -Name Name -Password Password -WindowsLiveID WindowsLiveId  
[-EvictLiveId <$true | $false>] {AddtlParams} {CommonParams}  
{ModParams}
```

```
New-Mailbox -Name Name -UseExistingLiveId <$true | $false> -WindowsLiveID  
WindowsLiveId [-BypassLiveId <$true | $false>] [-NetID NetID ]  
{AddtlParams} {CommonParams} {ModParams}
```

```
New-Mailbox -Name Name -UserPrincipalName UserNameAndSuffix [-MailboxPlan  
MailboxPlanId ] {CommonParams} {ModParams}
```

```
New-Mailbox -Name Name -AccountDisabled <$true | $false> [-MailboxPlan  
MailboxPlanId ] [-Password Password ] [-UsageLocation Location ]  
[-UserPrincipalName UserNameAndSuffix ] {CommonParams} {ModParams}
```

```
New-Mailbox -Name Name -ImportLiveId <$true | $false> -WindowsLiveID  
WindowsLiveId {AddtlParams} {CommonParams} {ModParams}
```

New-Mailbox -Name **Name** -RemovedMailbox **RemovedMailboxId** [-MailboxPlan **MailboxPlanId**] [-Password **Password**] {CommonParams} {ModParams}

New-Mailbox -Name **Name** -FederatedIdentity **FederatedId** -WindowsLiveID **WindowsLiveId** [-EvictLiveId <\$true | \$false>] [-NetID **NetID**]
{AddtlParams} {CommonParams}

New-Mailbox -Name **Name** -FederatedIdentity **FederatedId**
-MicrosoftOnlineServicesID **WindowsLiveId** [-NetID **NetID**]
{AddtlParams} {CommonParams}

New-Mailbox -Name **Name** -ArchiveDomain **SmtptDomain** -Password **Password**
-UserPrincipalName **UserNameAndSuffix** [-MailboxPlan **MailboxPlanId**]
[-RemoteArchive <\$true | \$false>] [-RemovedMailbox **RemovedMailboxId**]
{CommonParams} {ModParams}

New-Mailbox -Name **Name** -MicrosoftOnlineServicesID **WindowsLiveId** -Password **Password** {AddtlParams} {CommonParams} {ModParams}

New-Mailbox -Name **Name** [-UserPrincipalName **UserNameAndSuffix**] {CommonParams}
{ModParams}

New-Mailbox -Name **Name** -LinkedDomainController **DCName** -LinkedMasterAccount **Identity** [-LinkedCredential **Credential**] [-UserPrincipalName **UserNameAndSuffix**]
{CommonParams} {ModParams}

New-Mailbox -Name **Name** -Equipment <\$true | \$false>[-Password **Password**]
[-UserPrincipalName **UserNameAndSuffix**] {CommonParams} {ModParams}

New-Mailbox -Name **Name** -Shared <\$true | \$false>[-Password **Password**]
[-UserPrincipalName **UserNameAndSuffix**] {CommonParams} {ModParams}

New-Mailbox -Name **Name** [-Password **Password**] [-UserPrincipalName **UserNameAndSuffix**] {CommonParams} {ModParams}

New-Mailbox -Name **Name** -Arbitration <\$true | \$false> -UserPrincipalName **UserNameAndSuffix** [-Password **Password**] {CommonParams}

New-Mailbox -Name **Name** [-Password **Password**] [-UserPrincipalName **UserNameAndSuffix**] {CommonParams} {ModParams}

New-Mailbox -Name **Name** -Discovery <\$true | \$false>[-Password **Password**]
[-UserPrincipalName **UserNameAndSuffix**] {CommonParams}

New-Mailbox -Name **Name** -EnableRoomMailboxAccount <\$true | \$false>
-Room <\$true | \$false>[-MicrosoftOnlineServicesID **WindowsLiveId**]

[-RoomMailboxPassword Password] [-UserPrincipalName UserNameAndSuffix]
{CommonParams}

New-Mailbox -Name Name -PublicFolder <\$true | \$false> [-HoldForMigration
<\$true | \$false>] [-IsExcludedFromServingHierarchy <\$true | \$false>] {CommonParams}

{AddtlParams}
[-MailboxPlan PlanID] [-RemovedMailbox RemovedMailboxId]
[-UsageLocation Location]

{ModParams}
[-ArbitrationMailbox ModeratorMailbox] [-ModeratedBy Moderators]
[-ModerationEnabled <\$true | \$false>] [-SendModerationNotifications
<Never | Internal | Always>]

{CommonParams}
[-ActiveSyncMailboxPolicy MailboxPolicyId] [-AddressBookPolicy ABPolicyId]
[-Alias ExchangeAlias] [-Archive {\$true | \$false}] [-ArchiveDatabase
DatabaseId] [-Database DatabaseId] [-DisplayName Name]
[-DomainController FullyQualifiedName] [-ExternalDirectoryObjectID ObjectID]
[-FirstName FirstName] [-ImmutableId Id] [-Initials Initials] [-LastName
LastName] [-ManagedFolderMailboxPolicy MailboxPolicyId]
[-ManagedFolderMailboxPolicyAllowed {\$true | \$false}]
[-Organization OrgName] [-OrganizationalUnit OUName]
[-OverrideRecipientQuotas {\$true | \$false}] [-PrimarySmtpAddress
SmtpAddress] [-QueryBaseDNRestrictionEnabled <\$true | \$false>]
[-RemoteAccountPolicy PolicyId] [-RemotePowershellEnabled
<\$true | \$false>] [-ResetPasswordOnNextLogon <\$true | \$false>]
[-RetentionPolicy PolicyId] [-RoleAssignmentPolicy PolicyId]
[-SamAccountName PreWin2000Name] [-SharingPolicy PolicyId]
[-TargetAllMDBs <\$true | \$false>] [-ThrottlingPolicy PolicyId]

Usage

New-Mailbox -Name "Shane S. Kim" -Alias "shaneK"
-OrganizationalUnit "imaginedlands.local/Engineering"
-Database "Engineering Primary"
-UserPrincipalName "shaneK@imaginedlands.local" -SamAccountName "shaneK"
-FirstName "Shane" -Initials "S" -LastName "Kim"
-ResetPasswordOnNextLogon \$true -Archive \$true

Creating Online User Accounts with Mailboxes

You can create user accounts with mailboxes in Exchange Online. These accounts are then available in the online organization.

To create an online user account, follow these steps:

1. In Office 365 Admin Center, select **Users** in the Navigation menu and then select **Active Users** .

2. On the Active Users page, click **Add A User** . This opens the New User window, shown in Figure 4-6.
3. Type the user's first name and last name in the text boxes provided. These values are used to create the Display Name entry.
4. The Display Name and Name properties can't exceed 64 characters. As necessary, make changes to the Display Name. For example, you might want to type the name in LastName FirstName format or in FirstName LastName format.
5. In the User Name text box, type the user's logon name. Use the drop-down list to select the domain with which you want to associate the account. This sets the fully qualified logon name, such as mikejackson@imaginedlands.onmicrosoft.com (which is referred to as the logon ID with Exchange Online).

The screenshot shows a 'New User' window with a dark header bar. On the left is a circular profile picture placeholder with the initials 'MJ'. To its right, the name 'Mike Jackson' and email 'mikejackson@williamstane.com' are displayed. Below the header is a white form area with several input fields:

- 'First name' text box containing 'Mike'
- 'Last name' text box containing 'Jackson'
- 'Display name *' text box containing 'Mike Jackson'
- 'User name *' text box containing 'mikejackson'
- 'Domain' dropdown menu showing 'williamstane.com'
- 'Location' dropdown menu showing 'United States'

 A close button (X) is in the top right corner of the window.

FIGURE 4-6 Providing the details for the new user.

6. Expand the Contact Information panel by clicking on it and add contact information, such as Job Title and Department, as appropriate.
7. By default, Exchange Online will generate a temporary password for the user and email the account information along with the password to the email address associated with your logon. Exchange Online will also make the user change the password when the first sign in. To change these settings, expand the Password panel by clicking on it and then specify the desired options. You can create a password for the user, specify an alternative address for emailing the account information or both.
8. By default, the user account is created with no administrator access. If you are creating an account for an administrator, expand the Roles panel by clicking on it and then specify the desired options. You can specify that the user is a global administrator or create a custom administrator role.
9. Expand the Product Licenses panel by clicking on it and select a product plan, license or both to assign to the user. Click **Save** to create the user account and mailbox.

IMPORTANT The available licenses will depend on the license types previously purchased for your organization. If you don't have available subscriptions, the appropriate subscriptions and licenses will be purchased for you

automatically.

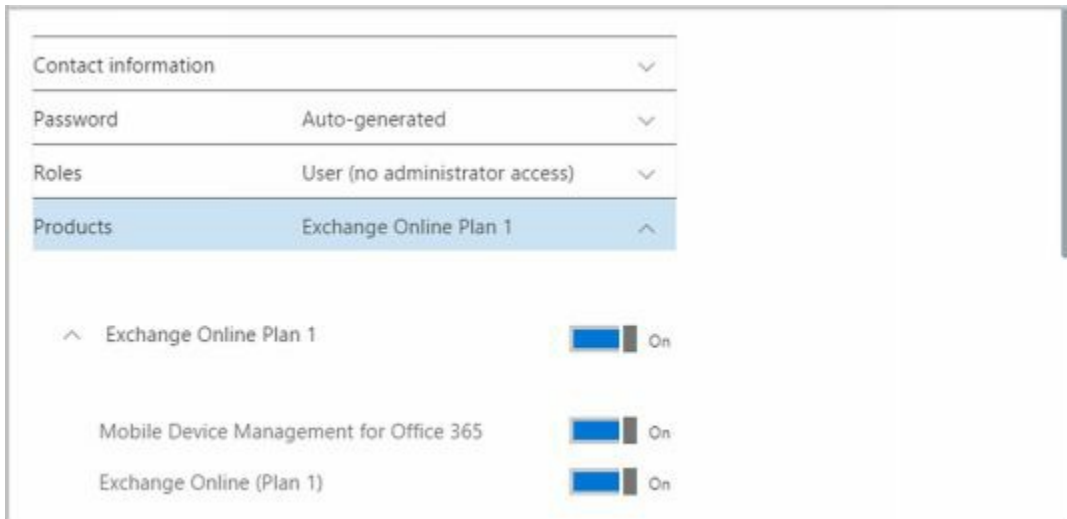


FIGURE 4-7 Providing additional account information and selecting product plans.

Creating the online user account and mailbox isn't necessarily the final step. You might also want to do the following:


- Add the user to security and distribution groups
- Enable or disable mailbox features for the account
- Modify the user's default delivery options, storage limits, and restrictions on the account
- Associate additional email addresses with the account

In Exchange Management Shell, you can create an online user account using the New-Mailbox cmdlet. Keep in mind that a mailbox is created only when you use the -MailboxPlan parameter to assign a mailbox plan to the new user.

Adding Mailboxes to Existing Domain User Accounts

You don't have to create an Exchange mailbox when you create a domain user account. You can create a mailbox for a domain user account any time you determine the mailbox is needed.

You can add a mailbox to an existing domain user account in several ways. If you are using a hybrid configuration and want the mailbox created in Exchange online, you can use the techniques discussed earlier under "Understanding on-premises and online recipient management." Otherwise, you can add a mailbox to a domain user account using only your on-premises Exchange administration tools. To do this, complete the following steps:

1. In Exchange Admin Center, select **Recipients** in the Navigation menu and then select **Mailboxes**.
2. Click **New** () and then select **User Mailbox**. This opens the New User Mailbox dialog box, shown in Figure 4-8.

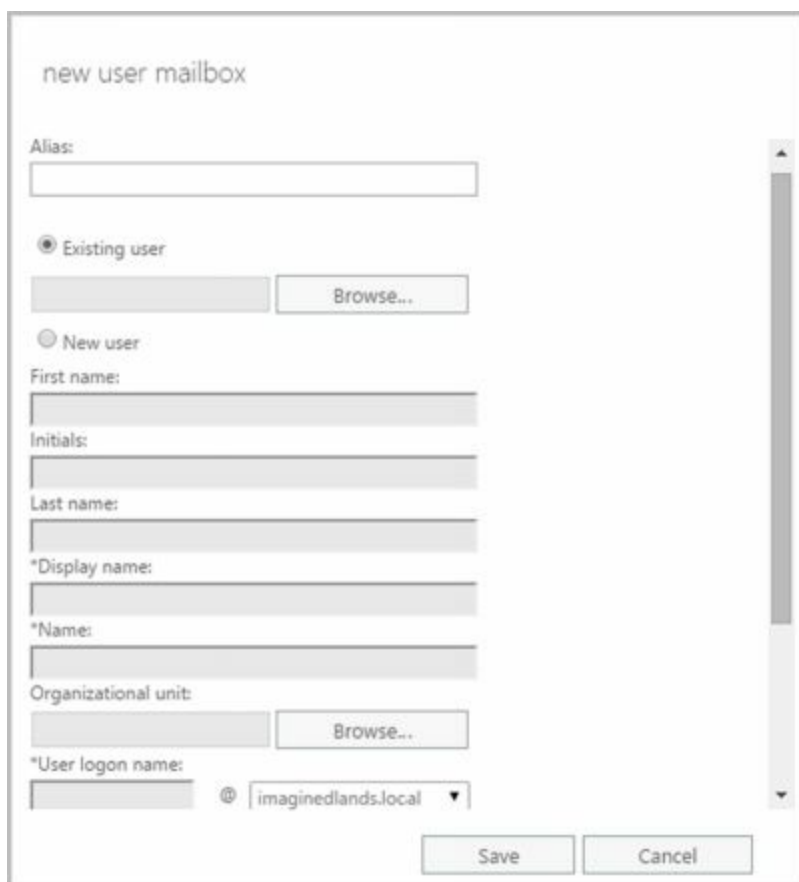


FIGURE 4-8 Adding a mailbox to an existing domain user account.

3. In the Alias text box, type an alias for the mailbox user. This alias should uniquely identify the user in the Exchange organization. Alias names cannot contain spaces.

NOTE The alias and domain suffix are combined to create the email address for the user. For example, if the alias is tedc and the domain suffix is imaginedlands.com, the email address is set as tedc@imaginedlands.com.

4. The Existing User option is selected by default. Click **Browse** . This displays the Select User dialog box.
5. In the Select User dialog box, shown in Figure 4-9, select the user account you want to mailbox-enable and then click **OK** . User accounts that are not yet mailbox-enabled or mailbox-enabled for the current domain are listed by name and organizational unit.

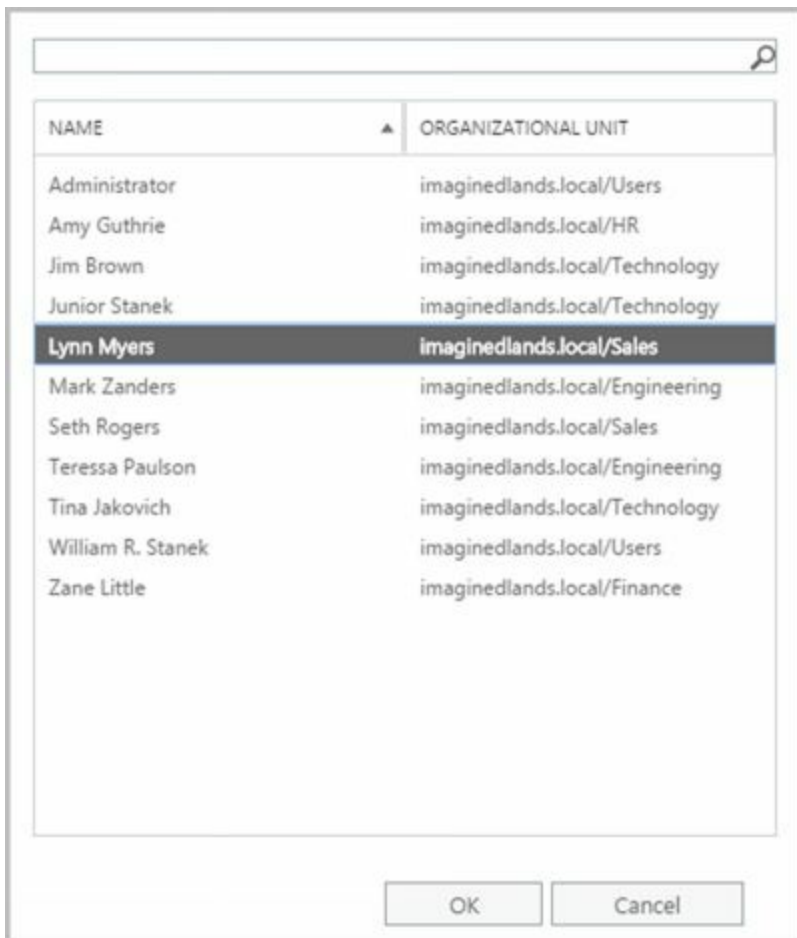


FIGURE 4-9 Finding the user account you want to mailbox-enable.

6. Click More Options. You can now:
 - **Specify the mailbox database** Exchange uses the mailbox provisioning load balancer to select a database to use when you create a mailbox and do not specify the mailbox database to use. If you want to specify the database to use, click Browse to the right of the Mailbox Database box. In the Select Mailbox Database dialog box, you'll see a list of available mailbox databases listed by name, server, and Exchange version. Select the mailbox database to use and then click OK.
 - **Create an archive mailbox** If you want to create an archive mailbox for the user, select the related check box. Items in the user's mailbox will be moved automatically to the archive mailbox based on the default retention policy. You also can choose a mailbox database for the archive. If you don't choose a mailbox database for the

archive, Exchange chooses one for you.

- **Assign an address book policy** By default, a user has access to the full address book information in the organization. Using address book policies, you can create customized address books. To apply an available policy, select it from the drop-down list.

7. Click **Save** . Exchange Admin Center creates the mailbox for the selected user. If an error occurs, the mailbox will not be created. You will need to click OK, correct the problem, and then click Save again.

In Exchange Management Shell, you can add a mailbox to individual user accounts using the Enable-Mailbox cmdlet. Sample 4-7 provides the syntax and usage. If you want to create mailboxes for multiple accounts, you need to enter a separate command for each account.

SAMPLE 4-7 Enable-Mailbox cmdlet syntax and usage

Syntax

```
Enable-Mailbox [-AccountDisabled <$true | $false>] [-MailboxPlan  
MailboxPlanId ] [-UsageLocation Location ] {AddtlParams} {CommonParams}
```

```
Enable-Mailbox -LinkedDomainController DCName -LinkedMasterAccount Identity [-  
Database DatabaseId ] [-LinkedCredential Credential ]  
[-TargetAllMDBs <$true | $false>] {CommonParams}
```

```
Enable-Mailbox -Discovery <$true | $false>[-Database DatabaseId ]  
[-TargetAllMDBs <$true | $false>] {CommonParams}
```

```
Enable-Mailbox [-AccountDisabled <$true | $false>] [-MailboxPlan  
MailboxPlanId ] [-UsageLocation Location ] {AddtlParams} {CommonParams}
```

```
Enable-Mailbox -Equipment <$true | $false> [-AccountDisabled <$true |  
$false>] {AddtlParams} {CommonParams}
```

```
Enable-Mailbox -Room <$true | $false> [-AccountDisabled <$true | $false>]  
{AddtlParams} {CommonParams}
```

```
Enable-Mailbox -PublicFolder <$true | $false>[-Database DatabaseId ]  
[-HoldForMigration <$true | $false>] {CommonParams}
```

```
Enable-Mailbox -Arbitration <$true | $false>[-Database DatabaseId ]  
[-TargetAllMDBs <$true | $false>] {CommonParams}
```

```
Enable-Mailbox -Shared <$true | $false> [-AccountDisabled <$true | $false>]  
{AddtlParams} {CommonParams}
```

```
Enable-Mailbox [-Archive <$true | $false>] [-ArchiveDatabase DatabaseId ]  
[-ArchiveGuid <Guid>] [-ArchiveName <MultiValuedProperty>]
```

[-BypassModerationCheck <\$true | \$false>] {CommonParams}

Enable-Mailbox -ArchiveDomain **SmtpDomain** [-RemoteArchive <\$true | \$false>] {CommonParams}

{AddtlParams}

[-BypassModerationCheck <\$true | \$false>] [-Database DatabaseId] [-TargetAllMDBs <\$true | \$false>]

{CommonParams}

[-ActiveSyncMailboxPolicy **MailboxPolicyId**] [-Alias **ExchangeAlias**]

[-DisplayName **Name**] [-DomainController **FullyQualifiedName**]

[-ManagedFolderMailboxPolicy **MailboxPolicyId**]

[-ManagedFolderMailboxPolicyAllowed {\$true | \$false}]

[-OverrideRecipientQuotas {\$true | \$false}]

[-PrimarySmtpAddress **SmtpAddress**]

[-RetentionPolicy **PolicyId**] [-RoleAssignmentPolicy **PolicyId**]

Usage

Enable-Mailbox -Identity "imaginedlands.local/Engineering/Oliver Lee"

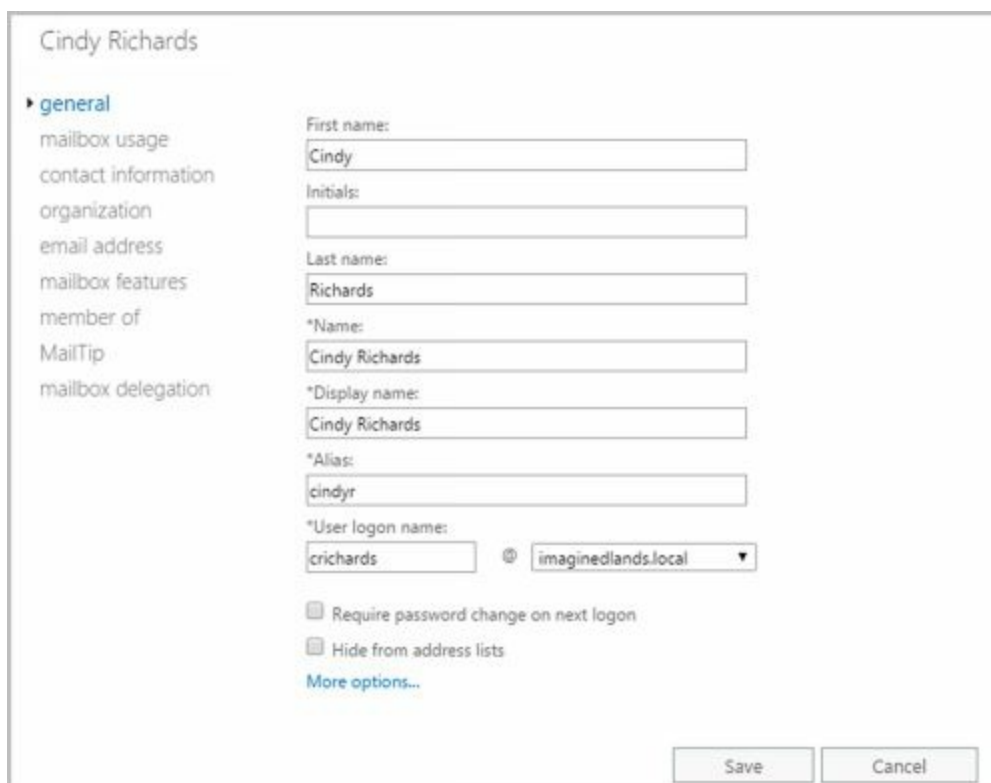
-Alias "Oliverl" -Database "Engineering Primary"

Setting or Changing the Common Name and Logon Name for Domain User Accounts

All domain user accounts have a common name stored in Active Directory and a logon name used for logging on to the domain. These names can be different from the mailbox display name and mailbox alias used by Exchange Server.

You can set this information for a domain user account by completing the following steps:

1. In Exchange Admin Center, select **Recipients** in the Navigation menu and then select **Mailboxes**.
2. Double-click the mailbox entry for the user with which you want to work. This opens a properties dialog box for the user.
3. On the General page, shown in Figure 4-10, use the following text boxes to set the user's common name and logon name:



The screenshot shows the 'Cindy Richards' user properties dialog box in the Exchange Admin Center. The 'general' tab is selected, showing various fields for user information. The fields and their values are:

- First name: Cindy
- Initials: (empty)
- Last name: Richards
- *Name: Cindy Richards
- *Display name: Cindy Richards
- *Alias: cindyr
- *User logon name: crichards @ imaginedlands.local

There are also two checkboxes: 'Require password change on next logon' and 'Hide from address lists', both of which are unchecked. A 'More options...' link is visible below the checkboxes. At the bottom right, there are 'Save' and 'Cancel' buttons.

FIGURE 4-10 Changing the user's naming information for Active Directory.

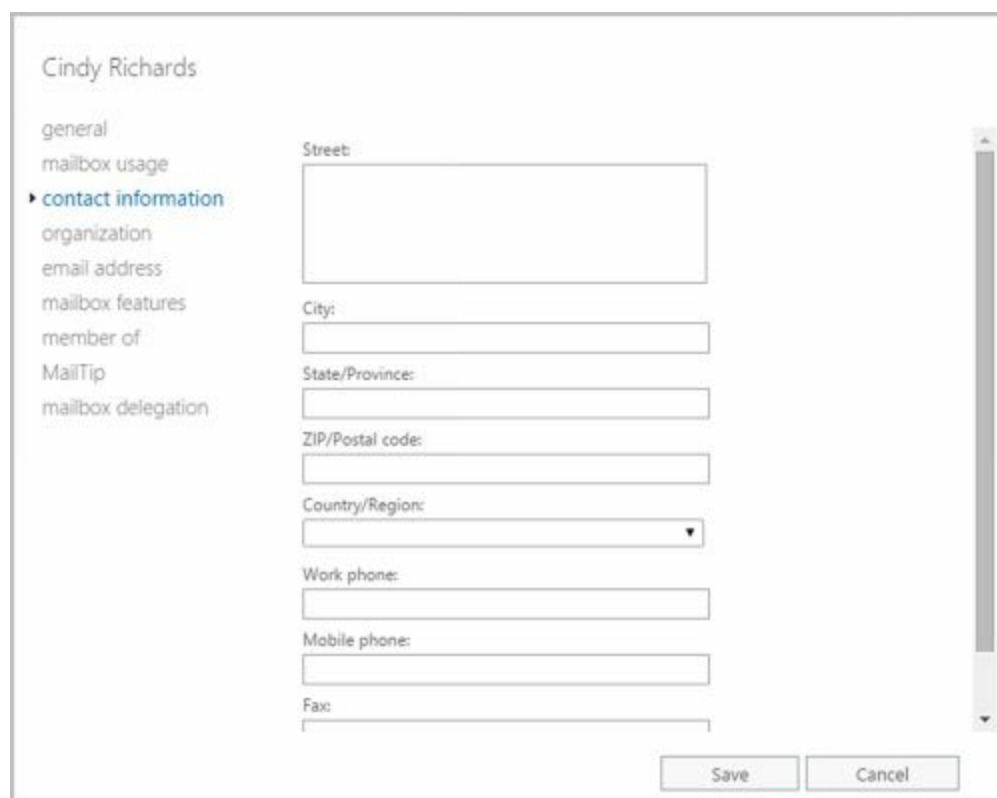
- **First Name, Initials, Last Name** Sets the user's full name.
 - **Name** Sets the user's display name as seen in logon sessions and in Active Directory.
 - **User Logon Name** Sets the user's logon name.
4. Click **Save** to apply your changes.

Setting or Changing Contact Information for User Accounts

You can set contact information for a user account by completing the following steps:

1. In Exchange Admin Center, select **Recipients** in the Navigation menu and then select **Mailboxes**.
2. Double-click the mailbox entry for the user with which you want to work.
3. On the Contact Information page, shown in Figure 4-11, use the text boxes provided to set the user's business address or home address. Normally, you'll want to enter the user's business address. This way, you can track the business locations and mailing addresses of users at various offices.

NOTE You need to consider privacy issues before entering private information, such as home addresses and home phone numbers, for users. Discuss the matter with the appropriate groups in your organization, such as the human resources and legal departments. You might also want to get user consent before releasing home addresses.



The screenshot shows the 'Contact Information' page for a user named Cindy Richards. The page is divided into a left sidebar and a main content area. The sidebar contains a list of navigation options: 'general', 'mailbox usage', 'contact information' (which is highlighted with a blue arrow), 'organization', 'email address', 'mailbox features', 'member of', 'MailTip', and 'mailbox delegation'. The main content area contains a form with the following fields: 'Street' (a large text box), 'City' (a text box), 'State/Province' (a text box), 'ZIP/Postal code' (a text box), 'Country/Region' (a dropdown menu), 'Work phone' (a text box), 'Mobile phone' (a text box), and 'Fax' (a text box). At the bottom right of the form, there are two buttons: 'Save' and 'Cancel'.

FIGURE 4-11 Setting contact information for a user.

4. Use the Work Phone, Mobile Phone, and Fax text boxes to set the user's primary business telephone, mobile phone, and fax numbers.
5. Click More Options. Use the Office text box to set the user's office and the Web Page text box to set the URL of the user's home page, which can be on the Internet or the company intranet.

6. On the Organization page, shown in Figure 4-12, as appropriate, type the user's title, department, and company.
7. To specify the user's manager, click **Browse** . In the Manager dialog box, select the user's manager and then click **OK** . When you specify a manager, the user shows up as a direct report in the manager's account. Click **Save** to apply the changes.

Cindy Richards

- general
- mailbox usage
- contact information
- ▶ **organization**
- email address
- mailbox features
- member of
- MailTip
- mailbox delegation

Title:

Department:

Company:

Manager:

Direct reports:

FIGURE 4-12 Adding organizational information for a user.

Changing Logon ID or Logon Domain for Online Users

For Exchange Online, the fully-qualified logon ID is the user's name followed by the @ symbol and the user's logon domain. You can modify this information for an online user account by completing the following steps:

1. In Office 365 Admin Center, select **Users** in the Navigation menu and then select **Active Users**.
2. Click the mailbox entry for the user with which you want to work. This opens a properties dialog box for the user.
3. Click **Edit** on the Email Addresses panel and then use the User Name and Domain text boxes to set the user's logon name and domain.
4. Click **Save** to apply your changes.

The screenshot shows the 'Edit email addresses' dialog box. It has the following sections:

- User name:** A text box containing 'wrstanek' and a 'Domain' dropdown menu showing 'imaginedlands.onmicrosoft.com'.
- Primary email address:** A text box containing 'wrstanek' and a 'Domain' dropdown menu showing 'imaginedlands.onmicrosoft.com'.
- Aliases:** A section with the text 'An alias is another email address where people can email William Stanek'. It contains an 'Alias' text box, a 'Domain' dropdown menu showing 'imaginedlands.onmicros', and a '+ Add' button.
- Email Address List:** A list showing 'wrstanek@imaginedlands.onmicrosoft.com' with a 'Set as primary' button and a delete icon.
- Buttons:** 'Save' and 'Cancel' buttons at the bottom.

FIGURE 4-13 Updating the user name for an Exchange Online user. Technet24.ir

Changing a User's Exchange Server Alias and Display Name

Each mailbox has an Exchange alias and display name associated with it. The Exchange alias is used with address lists as an alternative way of specifying the user in the To, Cc, or Bcc text boxes of an email message. The alias also sets the primary SMTP address associated with the account.

TIP Whenever you change the Exchange alias in the on-premises organization, a new email address is generated and set as the default address for SMTP. The previous email addresses for the account aren't deleted. Instead, these remain as alternatives to the defaults. To learn how to change or delete these additional email addresses, see "Adding, changing, and removing email and other addresses" later in this chapter.

With Exchange Online, changing a user's Exchange alias doesn't normally change the primary SMTP address for the user.

To change the Exchange alias and mailbox name on a user account, complete the following steps:

1. In Exchange Admin Center, select **Recipients** in the Navigation menu and then select **Mailboxes**.
2. Double-click the mailbox entry for the user with which you want to work.
3. On the General page, the Display Name text box sets the mailbox name. Change this text box if you'd like the mailbox to have a different display name.
4. The Alias text box sets the Exchange alias. If you'd like to assign a new alias, enter the new Exchange alias in this text box.
5. Click **Save**.

NOTE Often, the user logon name and the Exchange alias are set to the same value. If you've implemented this practice in your organization, you may also want to modify the user logon name. However, this is not a best practice when security is a concern.

Joe Montgomery

- general
- mailbox usage
- contact information
- organization
- email address
- mailbox features
- member of
- MailTip
- mailbox delegation

First name: Joe

Initials:

Last name: Montgomery

*Name: Joe Montgomery

*Display name: Joe Montgomery

*Alias: joem

*User logon name: joem @ imaginedlands.local

Require password change on next logon

Hide from address lists

[More options...](#)

Save Cancel

FIGURE 4-14 Updating the user name for an Exchange Server user.

Adding, Changing, and Removing Email and Other Addresses

When you create a mailbox-enabled user account, default email addresses are created. Any time you update the user's Exchange alias in the on-premises Exchange organization, a new default email address is created. However, the old addresses aren't deleted. They remain as alternative email addresses for the account.

With Exchange Online, changing a user's Exchange alias doesn't normally change the email address for the user. You can, however, modify the primary SMTP address or add additional SMTP addresses.


Exchange also allows you to create non-SMTP addresses for users:


- Exchange Unified Messaging (EUM) addresses used by the Unified Messaging service to locate UM-enabled users within the Exchange organization
- Custom addresses for legacy Exchange (Ex) as well as these non-Exchange mail organizations: X.400, X.500, MSMail, CcMail, Lotus Notes, and Novell GroupWise

To add, change, or remove an email or other address, follow these steps:

1. In Exchange Admin Center, select **Recipients** in the Navigation menu and then select **Mailboxes**.
2. Double-click the mailbox entry for the user you want to work with.
3. On the Email Address page, shown in Figure 4-15, you can use the following techniques to manage the user's email addresses:

- **Create a new SMTP address** Click Add (). Because the address type SMTP is selected by default, enter the SMTP email address, and then click OK to save your changes.

- **Create a new EUM address** Click Add (), and then select the EUM option. Enter the custom address or extension. Next, click Browse and then select a dial plan. Click OK to save your changes.

- **Create a custom address** Click Add (), and then select the Custom Address Type option. Enter the custom address type in the text box provided. Valid types include: X.400, X.500, EUM, MSMail, CcMail, Lotus Notes, and NovellGroupWise. Next, enter the custom address. This address must comply with the format requirements for the address type. Click OK to save your changes.

TIP Use SMTP as the address type for standard Internet email addresses. For custom address types, such as X.400, you must enter the address in the proper format.

- **Edit an existing address** Double-click the address entry, or select the entry and then select Edit on the toolbar. Modify the settings in the Address dialog box, and then click OK.
- **Delete an existing address** Select the address, and then click Remove.

NOTE You can't delete the primary SMTP address without first promoting another email address to the primary position. Exchange Server uses the primary SMTP address to send and receive messages.

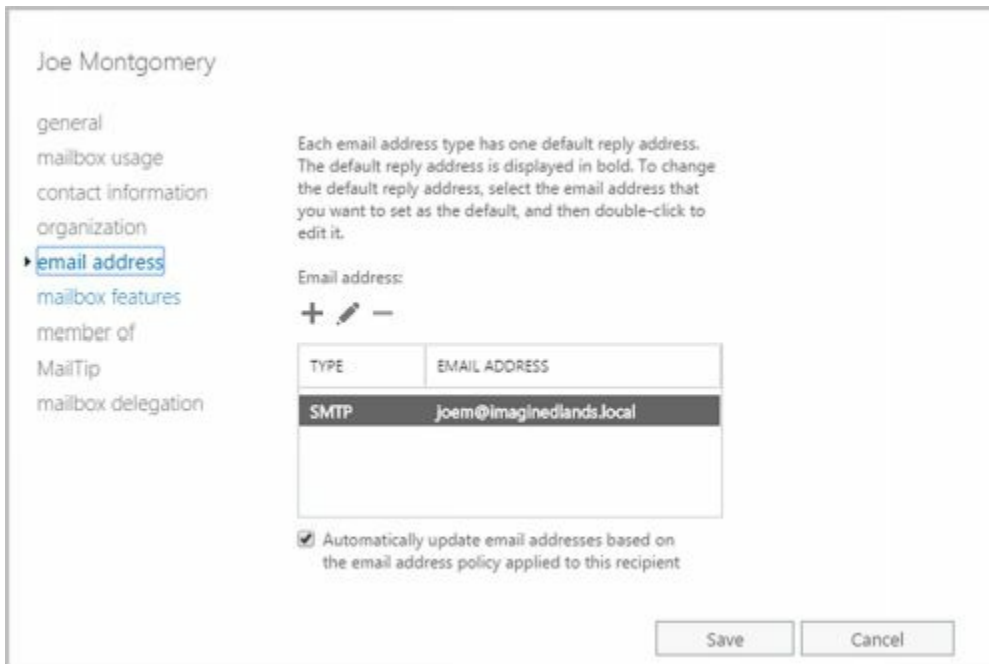



FIGURE 4-15 Configuring the email addresses for the user account.

Setting a Default Reply Address for a User Account

Each email address type has one default reply address. This email address sets the value of the Reply To text box. To change the default reply address, follow these steps:

1. In Exchange Admin Center, select **Recipients** in the Navigation menu and then select **Mailboxes** .
2. Double-click the mailbox entry for the user with which you want to work.
3. On the Email Address page, current default email addresses are highlighted with bold text. Email addresses that aren't highlighted are used only as alternative addresses for delivering messages to the current mailbox. To change the current default settings, select an email address that isn't highlighted and then click **Edit** ().
4. In the Email Address dialog box, select the **Make This The Reply Address** checkbox. Click **OK** to save the changes.

Changing A User's Web, Wireless Service, And Protocol Options

When you create user accounts with mailboxes, global settings determine the web, wireless services, and protocols that are available. You can change these settings for individual users at any time by completing the following steps:

1. In Exchange Admin Center, select **Recipients** in the Navigation menu and then select **Mailboxes**.
 2. Double-click the mailbox entry for the user with which you want to work.
 3. Click the **Mailbox Features** tab. As shown in Figure 4-16, configure the following web, wireless services, and protocols for the user:
- **Exchange ActiveSync** Allows the user to synchronize the mailbox and to browse wireless devices. Properties allow you to specify an Exchange ActiveSync policy. When you enable Exchange ActiveSync, the account uses the default mobile device mailbox policy. To set an alternative policy, click the related **View Details** option.
 - **Outlook Web App** Permits the user to access the mailbox with a web browser. Properties allow you to specify an Outlook Web App mailbox policy.
 - **Unified Messaging** Allows the user to access unified messaging features, such as the voice browser. In a standard configuration of Exchange 2016, all new mailbox users have unified messaging enabled. However, a default UM Mailbox policy is required to fully activate the feature. If one hasn't been assigned, click **Enable** to display a dialog box that will allow you to specify the required policy.
 - **MAPI** Permits the user to access the mailbox with a Messaging Application Programming Interface (MAPI) email client.

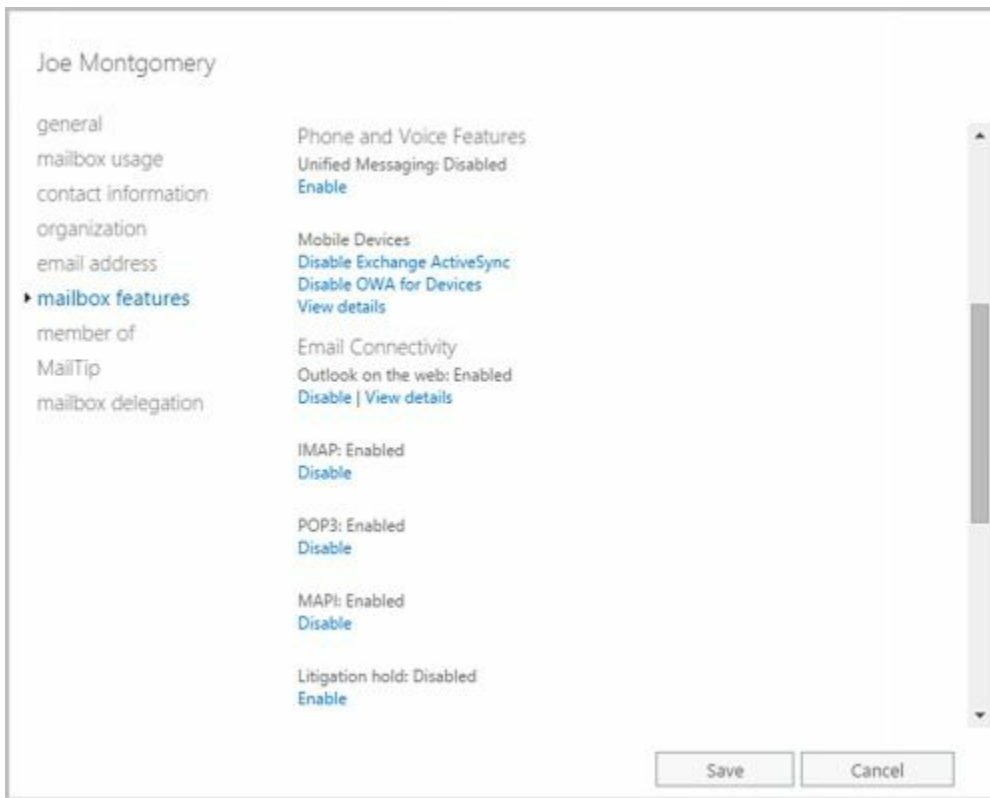


FIGURE 4-16 Changing mailbox options for users.

- **POP3** Permits the user to access the mailbox with a Post Office Protocol version 3 (POP3) email client.
 - **IMAP4** Permits the user to access the mailbox with an Internet Message Access Protocol version 4 (IMAP4) email client.
 - **Litigation Hold** Indicates whether a mailbox is subject to litigation hold where users can delete mail items but the items are retained by Exchange. Properties allow you to provide a note to users about litigation hold and the URL of a webpage where they can learn more.
 - **Archive** Indicates whether an in-place archive mailbox has been created for the user. When you enable an in-place archive, you can specify the mailbox database to use. Properties allow you to specify the name of the folder in the user's mailbox that contains the archive as well as to set archive quota limit and warning values.
4. Select an option and then click **Enable** or **Disable**, as appropriate, to change the status. If an option has required properties, you'll be prompted to configure these properties when you enable the option. If an option has additional configurable properties, click the related **View Details** option to configure them.
 5. Click **Save** to close the Properties dialog box.

Requiring Domain User Accounts to Change Passwords

Group Policy settings typically require users to periodically change their passwords. Sometimes, you might have to ensure that a user changes her password the next time she logs on. For example, if you have to reset a password and give it to the user over the phone, you might want the user to change the password the next time she logs on.

You can set a user account to require the password to be changed on next logon by completing the following steps:

1. In Exchange Admin Center, select **Recipients** in the Navigation menu and then select **Mailboxes**.
2. Double-click the mailbox entry for the user with which you want to work.
3. On the General page, select the **Require Password Change OnNext Logon** check box. Click **OK**.

You can use the Set-User cmdlet to perform the same task, following the syntax shown in Sample 4-8.

SAMPLE 4-8 Requiring a user password change

Syntax

```
Set-User -Identity UserIdentity  
-ResetPasswordOnNextLogon <{$false|$true}>
```

Usage

```
Set-User -Identity "Oliver Lee" -ResetPasswordOnNextLogon $true
```


Deleting Mailboxes from User Accounts

When you disable a mailbox for a domain user account using the Exchange management tools, you permanently remove all Exchange attributes from the user object in Active Directory and mark the primary mailbox for deletion. Exchange Server then deletes the mailbox according to the retention period you set on the account or on the mailbox database. Because you only removed the user account's Exchange attributes, the user account still exists in Active Directory.

In Exchange Admin Center, you can delete a mailbox from a domain user account and all related Exchange attributes by completing the following steps:

1. In Exchange Admin Center, select **Recipients** in the Navigation menu and then select **Mailboxes**.
2. Select the mailbox entry for the user with which you want to work with.
3. Select the **More** button (**⋮**) and then select **Disable**.
4. When prompted to confirm this action, select **Yes**. The mailbox is then in the disconnected state and will be removed when the retention period expires. If the account was subject to litigation hold, mail items subject to litigation hold are preserved as recoverable items until the litigation hold period expires.

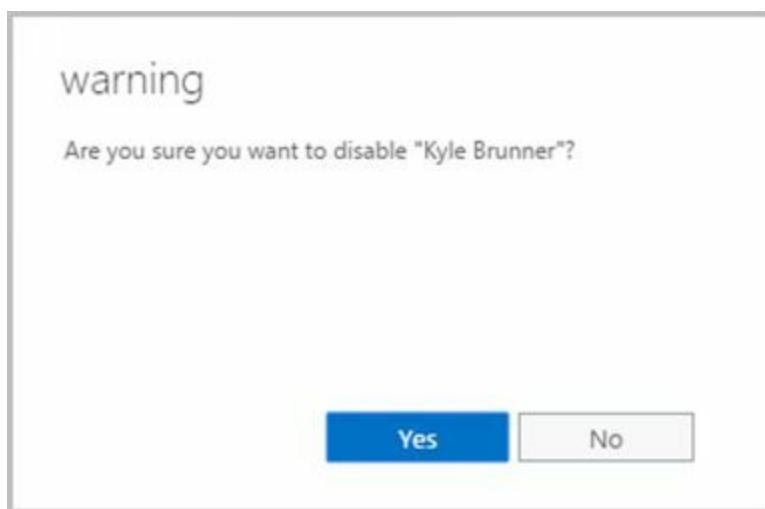


FIGURE 4-17 Disabling a mailbox and marking it for deletion.

If you remove the Exchange Online license for an online user account, the user's account is marked as an unlicensed account. Exchange Online deletes mailboxes from unlicensed accounts automatically after the grace period expires. By default, this grace period is 30 days. As with on-premises Exchange, retention hold, archiving and litigation hold settings determine whether some or any mailbox data is held.

You can remove a license from an online user account by completing the following steps:

1. In Office 365 Admin Center, select **Users** in the Navigation menu and then select **Active Users**.

2. Next, click the user whose license you want to remove.
3. Click **Edit** on the Product Licenses panel and then click the toggle to Off for the plan or license you want to remove.
4. Click **Assign**. The license that was previously assigned to this user will become available to be assigned to another user.

You can use the Disable-Mailbox cmdlet to delete mailboxes while retaining the user accounts as well. Sample 4-9 shows the syntax and usage.

SAMPLE 4-9 Disable-Mailbox cmdlet syntax and usage

Syntax

Disable-Mailbox -Identity **Identifier** [-DomainController **DCName**]

Usage


Disable-Mailbox -Identity "Oliver Lee"

Deleting User Accounts and Their Mailboxes

When you delete a domain user account and its mailbox using the Exchange management tools, you permanently remove the account from Active Directory and mark the primary mailbox for deletion. Exchange Server then deletes the mailbox according to the retention period you set on the account or on the mailbox database. Further, if the account was subject to litigation hold, mail items subject to litigation hold are preserved as recoverable items until the litigation hold period expires.

After you delete an account, you can't create an account with the same name and have the account automatically retain the same permissions as the original account. This is because the SID for the new account won't match the SID for the old account. However, that doesn't mean that after you delete an account, you can never again create an account with that same name. For example, a person might leave the company only to return a short while later. You can create an account using the same naming convention as before, but you'll have to redefine the permissions for that account.

Because deleting built-in accounts could have far-reaching effects on the domain, Windows doesn't let you delete built-in user accounts. In Exchange Admin Center, you can remove other types of accounts and the mailboxes associated with those accounts by following these steps:

1. In Exchange Admin Center, select **Recipients** in the Navigation menu and then select **Mailboxes**.
2. Select the user to delete and then click **Delete** ().
3. When prompted to confirm this action, select **Yes**.

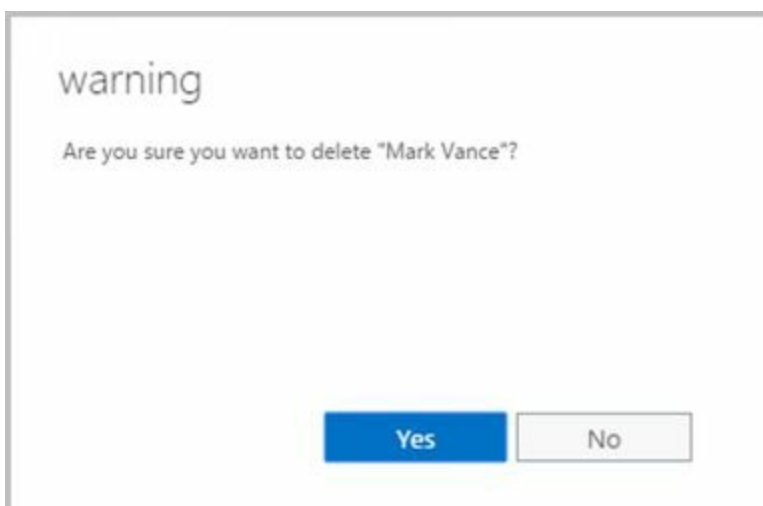


FIGURE 4-18 Confirming that you want to delete the account and mailbox.

NOTE Because Exchange security is based on domain authentication, you can't have a mailbox without an account. If you still need the mailbox for an account you want to delete, you can disable the account using Active Directory Users And Computers. Disabling the account in Active Directory prevents the user from

logging on, but you can still access the mailbox if you need to. To disable an account, Right-click the account in Active Directory Users And Computers and then select Disable Account. If you don't have permissions to use Active Directory Users And Computers, ask a domain administrator to disable the account for you.

IMPORTANT If your organization synchronizes user accounts to Exchange Online from your on-premises Active Directory environment, you must delete and restore synced user accounts using the on-premises tools. You can't delete or restore them in the online organization.

If you delete the corresponding Office 365 user account for a mailbox, the online user's mailbox is marked for deletion and the account is marked as a deleted account.

Deleted online users aren't removed immediately. Instead, the accounts are inactivated and marked for deletion. By default, the retention period is 30 days. When the retention period expires, a user and all related data is permanently deleted and is not recoverable. As with on-premises Exchange, retention hold, archiving, and litigation hold settings determine whether some or any mailbox data is held.

You can delete an online user account by completing the following steps:

1. In Office 365 Admin Center, select **Users** in the Navigation menu and then select **Active Users**.
2. Click the user whose license you want to remove and then click **Delete User**.
3. When prompted to confirm this action, select **Delete**. The license that was previously assigned to this user will become available to be assigned to another user.

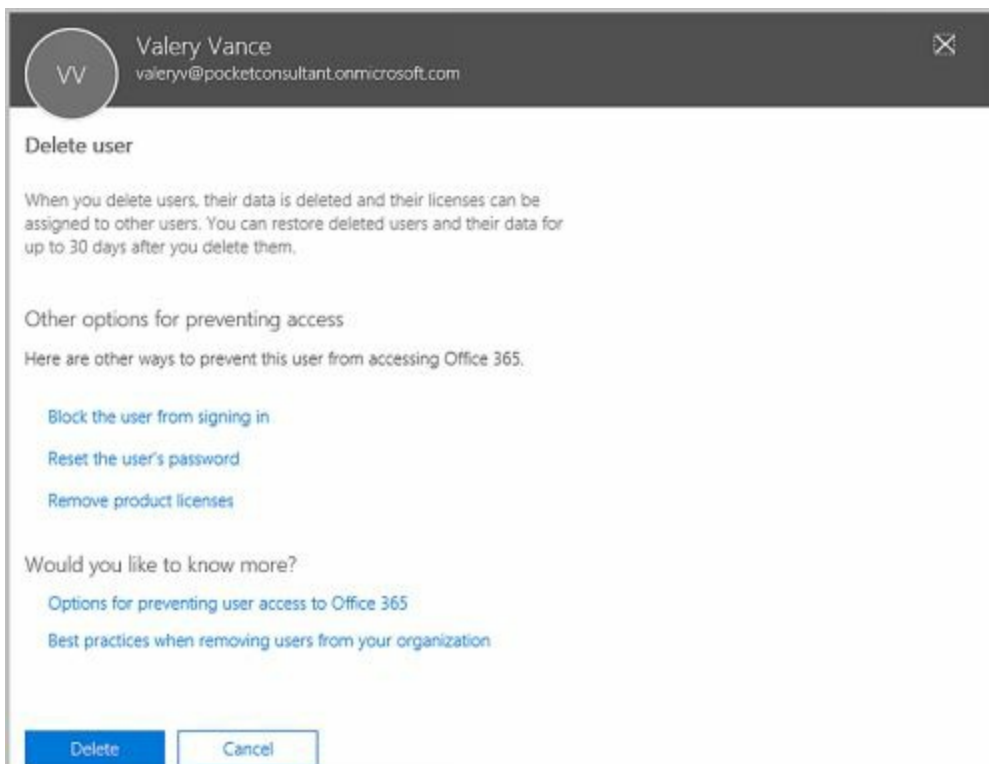


FIGURE 4-19 Confirming the deletion.

You also can use the Remove-Mailbox cmdlet to delete user accounts. Sample 4-10 shows the standard syntax. By default, the `-Permanent` flag is set to `$false` and mailboxes are retained in a disconnected state according to the mailbox retention policy. Otherwise, set the `-Permanent` flag to `$true` to remove the mailbox from Exchange.

SAMPLE 4-10 Remove-Mailbox cmdlet syntax and usage

Syntax

```
Remove-Mailbox -Identity UserIdentity {AddtlParams}
```

```
Remove-Mailbox -Database DatabaseId -StoreMailboxIdentity StoreMailboxId  
{AddtlParams}
```

```
{AddtlParams}
```

```
[-Arbitration <$false|$true>] [-DomainController DCName ]
```

```
[-IgnoreDefaultScope {$true | $false}] [-KeepWindowsLiveID {$true |  
$false}] [-Permanent <$false | $true>]
```

```
[-RemoveLastArbitrationMailboxAllowed {$true | $false}]
```

Usage

```
Remove-Mailbox -Identity "Oliver Lee"
```

```
Remove-Mailbox -Identity "Oliver Lee" -Permanent $true
```

Chapter 5. Managing Contacts

Contacts represent people with whom you or others in your organization want to get in touch. Contacts can have directory information associated with them, but they don't have network logon privileges.

The only difference between a standard contact and a mail-enabled contact is the presence of email addresses. A mail-enabled contact has one or more email addresses associated with it; a standard contact doesn't. When a contact has an email address, you can list the contact in the global address list or other address lists. This allows users to send messages to the contact.

In Exchange Admin Center, mail-enabled contacts and mail-enabled users are both listed in the Mail Contact node. Mail-enabled contacts are listed with the recipient type Mail Contact, and mail-enabled users are listed with the recipient type Mail User.

Creating Mail-Enabled Contacts

You can create and mail-enable a new contact by completing the following steps:

1. In Exchange Admin Center, select **Recipients** in the Navigation menu and then select **Contacts** .
2. Click **New** () and then select **Mail Contact** . This opens the New Mail Contact dialog box, shown in Figure 5-1.
3. Type the contact's first name, middle initial, and last name in the text boxes provided. These values are used to automatically create the following entries:
 - **Name** The common name is displayed in Active Directory (and only applies with on-premises Exchange).
 - **Display Name** The Display Name is displayed in the global address list and other address lists created for the organization. It is also used when addressing email messages to the contact.
4. Enter the Exchange alias for the contact. Aliases provide an alternative way of addressing users and contacts in To, Cc, and Bcc text boxes of email messages.
5. In the External Email Address text box, enter the address to associate with the contact. With on-premises Exchange, you can use both SMTP and non-SMTP addresses. With online Exchange, only standard SMTP addresses are accepted.

NOTE For non-SMTP addresses, the dialog box requires that you use a prefix that identifies the address type and that the address format comply to the rules for that type. Use the prefix X400: for X.400 addresses, the prefix X500: for X.500 addresses, the prefix MSMAIL: for MSMail addresses, the prefix CCMail: for CcMail addresses, the prefix LOTUSNOTES: for Lotus Notes, and the prefix NOVELLGROUPWISE: for NovellGroupWise.

FIGURE 5-1 Creating a new mail contact for the Exchange organization

6. The Organizational Unit text box shows where in Active Directory the contact will be created. By default, this is the Users container in the current domain. Because you'll usually need to create new contacts in a specific organizational unit rather than in the Users container, click **Browse** . Use the Select An Organizational Unit dialog box to choose the location in which to store the contact, and then click **OK** .
7. Click **Save** . Exchange Admin Center creates the new contact and mail-enables it. If an error occurs, the contact will not be created. You will need to correct the problem and repeat this procedure.

In Exchange Management Shell, you can create a new mail-enabled contact using the New-MailContact cmdlet. Sample 5-1 provides the syntax and usage.

SAMPLE 5-1 New-MailContact cmdlet syntax and usage

Syntax

```
New-MailContact -Name Name -ExternalEmailAddress TYPE:EmailAddress
[-ArbitrationMailbox ModeratorMailbox ] [-Alias ExchangeAlias]
[-DisplayName Name ] [-DomainController DCName ] [-FirstName FirstName ]
[-Initials Initials] [-LastName LastName ] [-MacAttachmentFormat <BinHex |
```


UuEncode | AppleSingle | AppleDouble>] [-MessageBodyFormat <Text | Html | TextAndHtml>] [-MessageFormat <Text | Mime>] [-ModeratedBy **Moderators**] [-ModerationEnabled <\$true | \$false>] [-Organization **OrgName**] [-OrganizationalUnit **OUName**] [-PrimarySmtpAddress **SmtpAddress**] [-SendModerationNotifications <Never | Internal | Always>] [-UsePreferMessageFormat <\$true | \$false>]

Usage

```
New-MailContact -ExternalEmailAddress "SMTP:wendywheeler@msn.com"
-Name "Wendy Wheeler" -Alias "WendyWheeler"
-OrganizationalUnit "imaginedlands.local/Corporate Services"
-FirstName "Wendy" -Initials "" -LastName "Wheeler"
```

In Exchange Management Shell, you can mail-enable an existing contact using the Enable-MailContact cmdlet. Sample 5-2 provides the syntax and usage.

SAMPLE 5-2 Enable-MailContact cmdlet syntax and usage

Syntax

```
Enable-MailContact -Identity ContactId -ExternalEmailAddress EmailAddress
[-Alias ExchangeAlias ] [-DisplayName Name ] [-DomainController
FullyQualifiedName ] [-MacAttachmentFormat <BinHex | UuEncode |
AppleSingle | AppleDouble>] [-MessageBodyFormat <Text | Html |
TextAndHtml>] [-MessageFormat <Text | Mime>] [-PrimarySmtpAddress
SmtpAddress ] [-UsePreferMessageFormat <$true | $false>]
```

Usage

```
Enable-MailContact -Identity "cpand.com/Sales/John Smith"
-ExternalEmailAddress "SMTP:johnsmith@imaginedlands.com"
-Alias "JohnSmith" -DisplayName "John Smith"
```

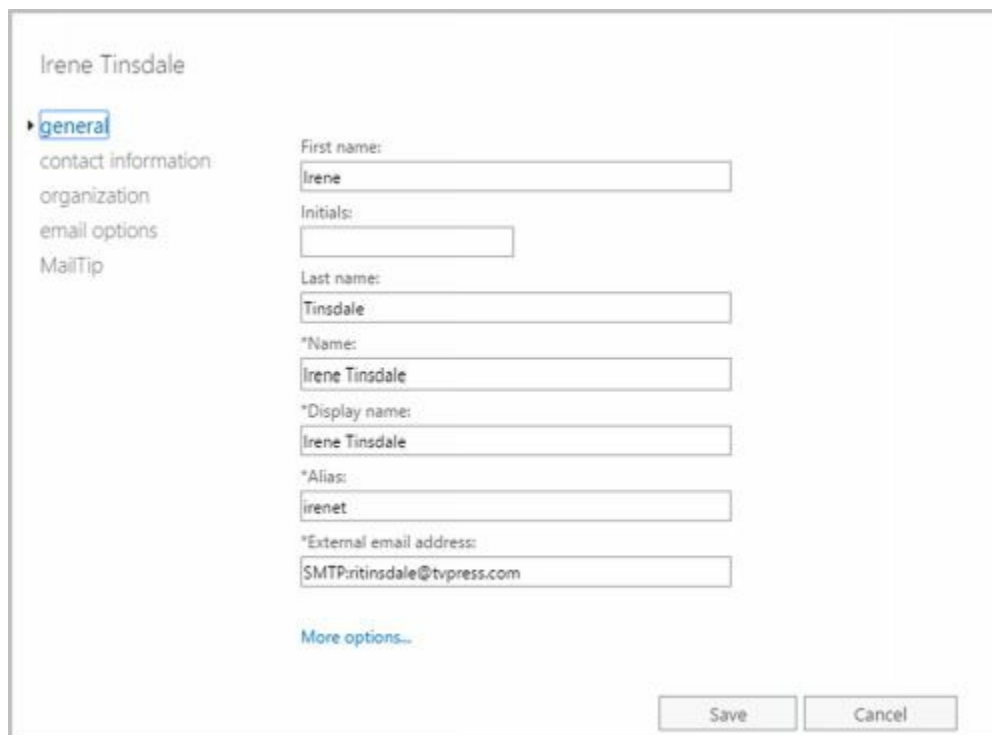
Setting or Changing a Contact's Name and Alias

Mail-enabled contacts can have the following name components:

- **First Name, Initials, Last Name** The first name, initials, and last name of the contact
- **CommonName** The name used in Active Directory, for on-premises contacts
- **Display Name** The name displayed in the global address list
- **Alias** The Exchange alias for the contact

You can set or change name and alias information for a mail-enabled contact or user by completing the following steps:

1. In Exchange Admin Center, select **Recipients** in the Navigation menu and then select **Contacts**.
2. Double-click the name of the mail-enabled contact or user you want to work with. The Properties dialog box appears.
3. On the General tab, use the textboxes provided to update the first name, middle initial, and last name as necessary. Changes you make will update the display name but not the common name. Therefore, as necessary, use the Name text box to update the common name.
4. With mail-enabled contacts, the Alias text box sets the Exchange alias. If you'd like to assign a new alias, enter the new Exchange alias in this text box.
5. With mail-enabled users, the User Logon Name text box sets the name used to log on to the domain as well as the domain suffix.
6. Click **Save** to apply your changes.



The screenshot shows the 'Irene Tinsdale' contact properties dialog box in Exchange Admin Center. The 'general' tab is selected, showing a sidebar with 'contact information', 'organization', 'email options', and 'MailTip'. The main area contains the following fields:

- First name: Irene
- Initials: (empty)
- Last name: Tinsdale
- *Name: Irene Tinsdale
- *Display name: Irene Tinsdale
- *Alias: irenet
- *External email address: SMTP:iritinsdale@tvpress.com

At the bottom, there is a 'More options...' link and 'Save' and 'Cancel' buttons.

FIGURE 5-2 Updating a contact.

Setting Additional Directory Information for Contacts

You can set additional directory information for a mail-enabled contact or user by completing the following steps:

1. In Exchange Admin Center, select **Recipients** in the Navigation menu and then select **Contacts**.
2. Double-click the name of the mail-enabled contact or user you want to work with. The Properties dialog box appears.
3. On the Contact Information page, use the text boxes provided to set the contact's business address or home address. Normally, you'll want to enter the contact's business address. This way, you can track the business locations and mailing addresses of contacts at various offices.

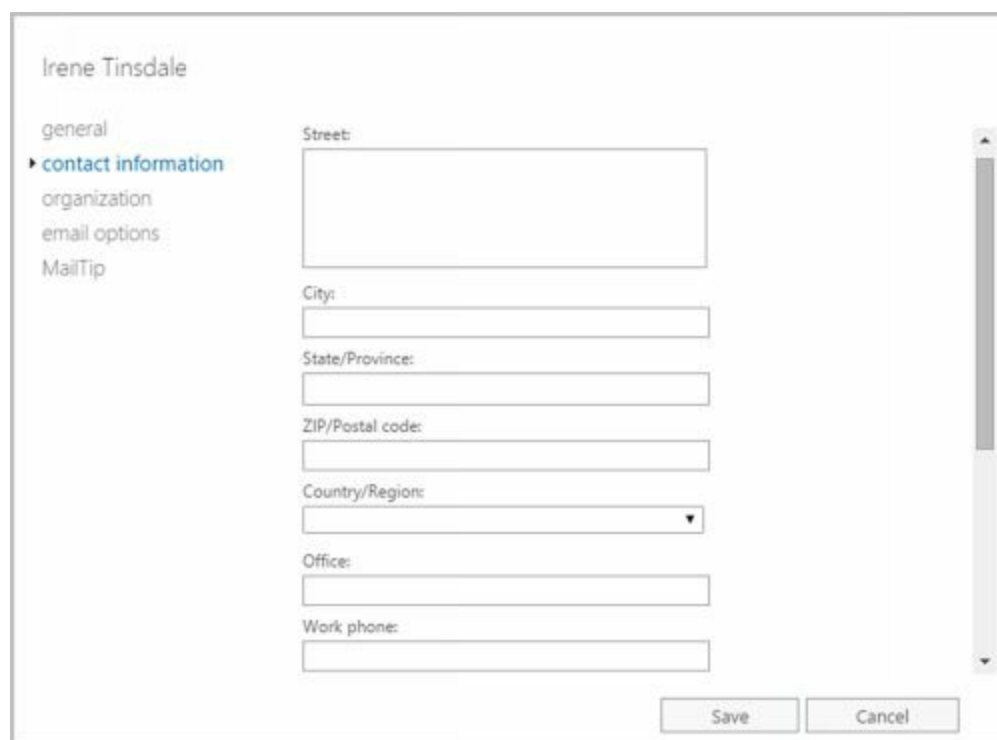
The screenshot shows the 'Contact Information' page for a contact named Irene Tinsdale. On the left, there is a navigation pane with options: 'general', 'contact information' (which is selected and highlighted), 'organization', 'email options', and 'MailTip'. The main area contains several text input fields and a dropdown menu for address information: 'Street:', 'City:', 'State/Province:', 'ZIP/Postal code:', 'Country/Region:' (with a dropdown arrow), 'Office:', and 'Work phone:'. At the bottom right of the form, there are two buttons: 'Save' and 'Cancel'.

FIGURE 5-3 Adding additional information to a contact.

NOTE You need to consider privacy issues before entering private information, such as home addresses and home phone numbers, for users. Discuss the matter with the appropriate groups in your organization, such as the human resources and legal departments. You might also want to get user consent before releasing home addresses.

4. Use the Work Phone, Mobile Phone, and Fax text boxes to set the contact or user's primary business telephone, mobile phone, and fax numbers.
5. Use the Office text box to set the user's Office and the Notes text box to add any important notes about the contact.

6. On the Organization page, as appropriate, type the contact or user's title, department, and company.
7. To specify the contact or user's manager, click **Browse** . In the Manager dialog box, select the manager and then click **OK** . When you specify a manager, the contact or user shows up as a direct report in the manager's account. Click **Save** to apply the changes.

Changing Email Addresses Associated with Contacts

Mail-enabled contacts and users have several types of email addresses associated with them:



- An internal, automatically generated email address used for routing within the organization
- An external email address to which mail routed internally is forwarded for delivery

With mail-enabled contacts, you can only use SMTP email addresses. You can change the SMTP email addresses associated with a mail-enabled contact by completing the following steps:

1. In Exchange Admin Center, select **Recipients** in the Navigation menu and then select **Contacts**.
2. Double-click the name of the mail-enabled contact you want to work with. The Properties dialog box appears.
3. On the General page, the external SMTP email address of the mail-enabled contact is listed. This is the primary SMTP email address for the mail-enabled contact. As necessary, enter a new email address.

NOTE The primary email address is listed with the prefix SMTP:. When you enter a new email address, you aren't required to enter this prefix. Thus, you could enter SMTP:williams@tresearch.net or williams@tresearch.net.

4. On the Email Options page, the primary SMTP email address is listed along with the internal email address. You can use the following techniques to manage the internal addresses:

- **Create an alternative internal address** Click Add (). Specify the internal email address to use by entering the Exchange alias and then selecting the domain for this internal address. Click OK.
 - **Edit an existing address** Double-click the address, or click Edit () on the toolbar. Modify the address settings as necessary, and then click OK.
 - **Delete an existing address** Select the address, and then click Remove.
5. Click **Save** to apply your changes.

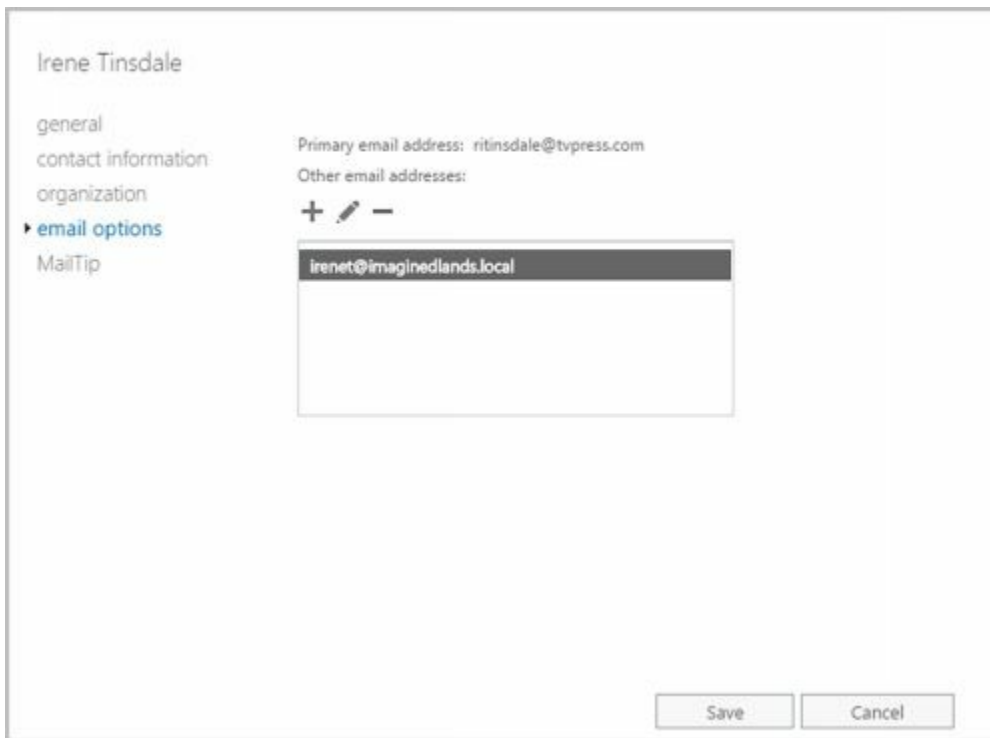



FIGURE 5-4 Modifying the email address information for a mail-enabled contact.

With mail-enabled users, you can use SMTP and non-SMTP email addresses. You can change the email addresses associated with a mail-enabled user by completing the following steps:

1. In Exchange Admin Center, select **Recipients** in the Navigation menu and then select **Contacts**.
2. Double-click the name of the mail-enabled user you want to work with. The Properties dialog box appears.
3. On the Email Addresses page, you can use the following techniques to manage the mail-enabled user's email addresses:
 - **Create a new SMTP address** Click Add (). Because the address type SMTP is selected by default, enter the SMTP email address, and then click OK to save your changes.

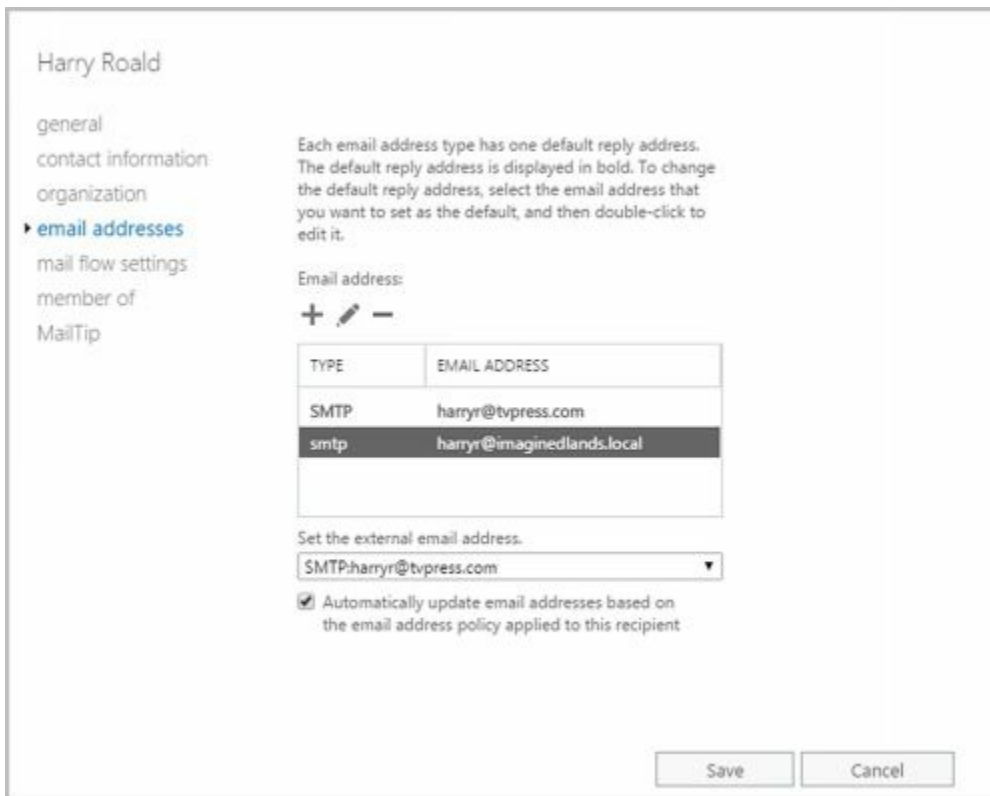



FIGURE 5-5 Modifying the email addresses for a mail-enabled user.

- **Create a custom address** Click Add (), and then select the Custom Address Type option. Enter the custom address type in the text box provided. Valid types include: X.400, X.500, EUM, MSMail, CcMail, Lotus Notes, and NovellGroupWise. Next, enter the custom address. This address must comply with the format requirements for the address type. Click OK to save your changes.
- **Edit an existing address** Double-click the address entry, or select the entry and then select Edit on the toolbar. Modify the settings in the Address dialog box, and then click OK.
- **Delete an existing address** Select the address, and then click Remove.

NOTE You can't delete the primary SMTP address without first promoting another email address to the primary position. Exchange Server uses the primary SMTP address to send and receive messages.

4. The external email address of the mail-enabled user is also listed on the Email Addresses page. This is the primary email address for the mail user or contact. As necessary, select an alternative email address to be the primary.
5. Click **Save** to apply your changes.

Disabling Contacts and Removing Exchange Attributes

With on-premises Exchange, you have two options for mail-enabled users and contacts that are no longer needed. You can disable the mail-enabled user or contact, or you can delete the mail-enabled user or contact. With Exchange online, your only option is to delete the mail-enabled user or contact.

When you disable a contact using the on-premises Exchange management tools, you permanently remove the contact from the Exchange database, but you do not remove it from Active Directory.

In Exchange Admin Center, you can disable mail-enabled contacts by following these steps:

1. In Exchange Admin Center, select **Recipients** in the Navigation menu and then select **Contacts**.
2. Select the contact that you want to disable.
3. Click the **More** button (**⋮**) and then select **Disable**.
4. When prompted to confirm this action, select **Yes**.

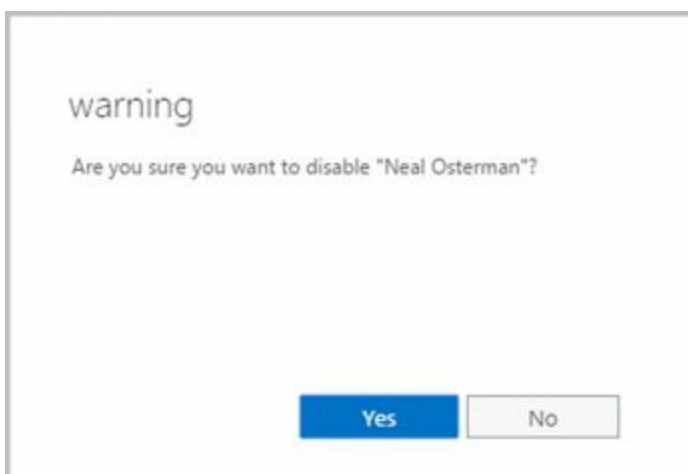


FIGURE 5-6 Disabling a contact.

You can use the `Disable-MailContact` cmdlet to remove Exchange attributes from contacts while retaining the contact in Active Directory. Sample 5-3 shows the syntax and usage.

SAMPLE 5-3 `Disable-MailContact` cmdlet syntax and usage

Syntax

```
Disable-MailContact -Identity ContactIdentity
```

Usage

```
Disable-MailContact -Identity "David So"
```


Later, if you want to re-enable the contact, you can do this using the `Enable-MailContact` cmdlet.

Deleting Contacts

When you delete a mail-enabled user or contact from Exchange Online, the mail-enabled user or contact is permanently removed from Exchange Online. When you delete a contact using the on-premises Exchange management tools, you permanently remove it from Active Directory and from the Exchange database. In Exchange Admin Center, you can delete contacts by following these steps:

1. In Exchange Admin Center, select **Recipients** in the Navigation menu and then select **Contacts** .
2. Select the contact that you want to delete and then click **Delete** .
3. When prompted to confirm this action, select **Yes** .

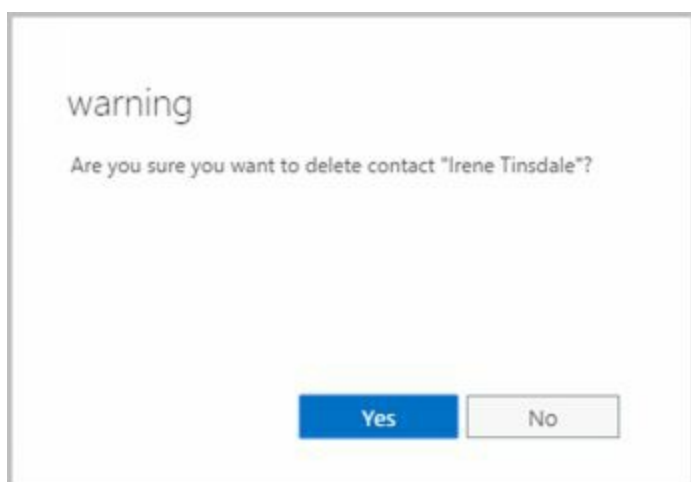


FIGURE 5-7 Deleting a contact.

You can use the `Remove-MailContact` cmdlet to delete contacts as well. Sample 5-4 shows the syntax and usage.

SAMPLE 5-4 Remove-MailContact cmdlet syntax and usage

Syntax

```
Remove-MailContact -Identity ContactIdentity
```

Usage

```
Remove-MailContact -Identity "Henrik Larsen"
```

Chapter 6. Adding Special-Purpose Mailboxes

Exchange Server 2016 and Exchange Online make it easy to create several special-purpose mailbox types, including:

- **Room mailbox** A room mailbox is a mailbox for room scheduling.
- **Equipment mailbox** An equipment mailbox is a mailbox for equipment scheduling.
- **Linked mailbox** A linked mailbox is a mailbox for a user from a separate, trusted forest.
- **Archive mailbox** An archive mailbox is used to store a user's messages, such as might be required for executives and needed by some managers.
- **Arbitration mailbox** An arbitration mailbox is used to manage approval requests, such as may be required for handling moderated recipients and distribution group membership approval.
- **Discovery mailbox** A Discovery mailbox is the target for Discovery searches and can't be converted to another mailbox type after it's created. In-Place eDiscovery is a feature of Exchange 2016 that allows authorized users to search mailboxes for specific types of content as might be required to meet legal discovery requirements.
- **Shared mailbox** A shared mailbox is a mailbox that is shared by multiple users, such as a general mailbox for customer inquiries.
- **Public folder mailbox** A public folder mailbox is a shared mailbox for storing public folder data.

The sections that follow discuss techniques for working with these special-purpose mailboxes.

Using Room and Equipment Mailboxes

You use room and equipment mailboxes for scheduling purposes only. You'll find that

- Room mailboxes are useful when you have conference rooms, training rooms, and other rooms for which you need to coordinate the use.
- Equipment mailboxes are useful when you have projectors, media carts, or other items of equipment for which you need to coordinate the use. Every room and equipment mailbox must have a separate user account associated with it. Although these accounts are required so that the mailboxes can be used for scheduling, the accounts are disabled by default so that they cannot be used for logon. To ensure that the resource accounts do not get enabled accidentally, you need to coordinate closely with other administrators in your organization.

IMPORTANT Each room or piece of equipment must have a separate user account. This is necessary to track the unique free/busy data for each resource.

NOTE The Exchange Admin Center doesn't show the enabled or disabled status of user accounts. The only way to check the status is to use domain administration tools.

Because the number of scheduled rooms and amount of equipment grows as your organization grows, you'll want to carefully consider the naming conventions you use with rooms and equipment:

- With rooms, you may want to use display names that clearly identify the rooms' physical locations. For example, you might have rooms named "Conference Room B on Fifth Floor" or "Building 83 Room 15."
- With equipment, you may want the display name to identify the type of equipment, the equipment's characteristics, and the equipment's relative location. For example, you might have equipment named "Dell LEDProjector at Seattle Office" or "Fifth Floor Media Cart."

As with standard user mailboxes, room and equipment mailboxes have contact information associated with them (see Figure 6-1). To make it easier to find rooms and equipment, you should provide as much information as possible. If a room has a conference or call-in phone, be sure to provide this phone number. Also, provide location details that help people find the conference room and specify the room capacity. The phone, location, and capacity are displayed in Office Outlook.

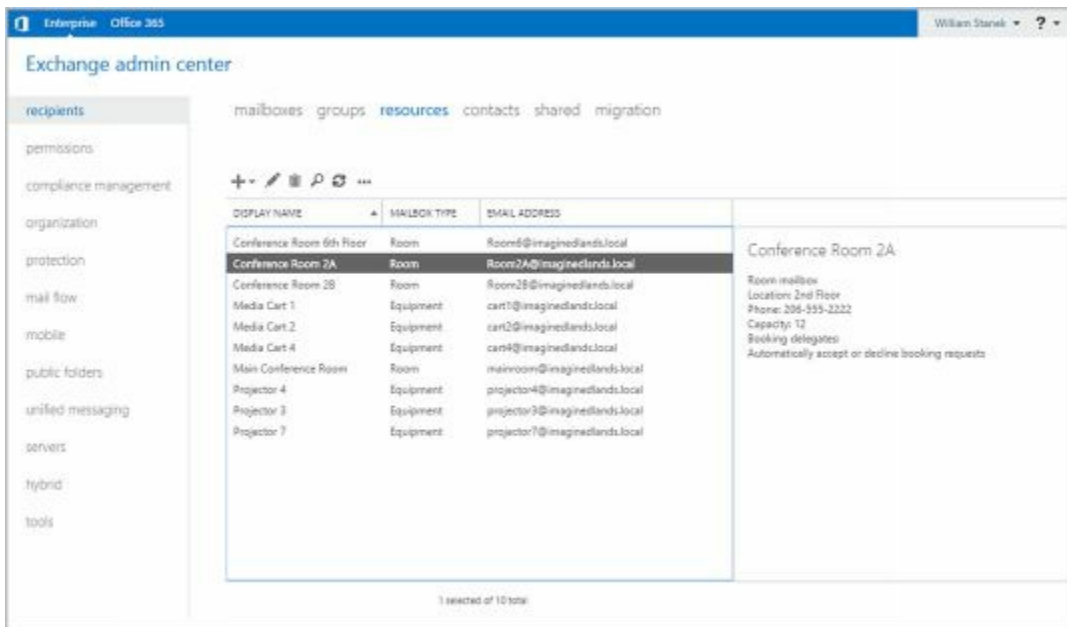


FIGURE 6-1 Mailboxes created for rooms and equipment.

After you've set up mailboxes for your rooms and equipment, scheduling the rooms and equipment is straightforward. In Exchange, room and equipment availability is tracked using free/busy data. In Outlook, a user who wants to reserve rooms, equipment, or both simply makes a meeting request that includes the rooms and equipment that are required for the meeting.

The steps to schedule a meeting and reserve equipment are as follows:

1. Create a meeting request. In Outlook 2010 or later, click **New Items**, and then select **Meeting**. Or press Ctrl+Shift+Q.
2. In the To text box, invite the individuals who should attend the meeting by typing their display names, Exchange aliases, or email addresses, as appropriate (see Figure 6-2).
3. Type the display name, Exchange alias, or email address for any equipment you need to reserve.
4. Click Rooms to the right of the Location text box. The Select Rooms dialog box appears, as shown in Figure 6-3. By default, the Select Rooms dialog box uses the All Rooms address book. Rooms are added to this address book automatically when you create them.




FIGURE 6-2 You can schedule a meeting with a reserved room and reserved equipment.

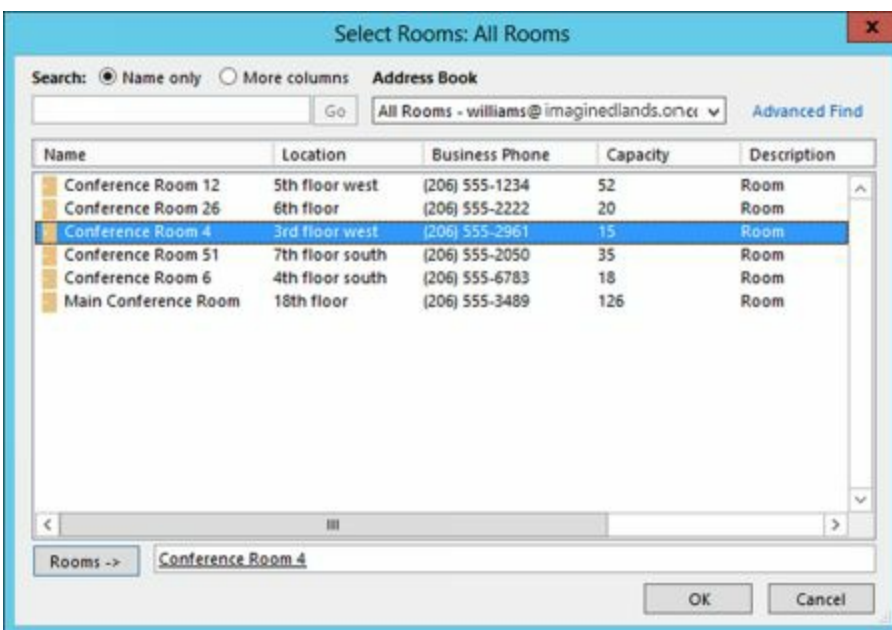


FIGURE 6-3 Select a room to use for the meeting.

5. Double-click the room you'd like to use. This adds the room to the Rooms list. Click **OK** to close the Select Rooms dialog box.
6. In the Subject text box, type the meeting subject.
7. Use the Start Time and End Time options to schedule the start and end times for the meeting.
8. Click **Scheduling Assistant** to view the free/busy data for the invited users and the selected resources. Use the free/busy data to make changes if necessary.
9. After you type a message to accompany the meeting request, click **Send**.

Exchange can be configured to accept booking requests automatically, based on availability, or to route requests through delegates, such as office administrators, who

review requests. Although small organizations might not need coordinators for rooms and equipment, most large organizations will need coordinators to prevent conflicts.

Both on-premises Exchange and Exchange Online provide additional booking options that can help to reduce conflicts (see Figure 6-4).

The screenshot shows the 'Conference Room 2B' dialog box in Outlook. The 'booking options' tab is selected. Under the heading 'Specify when this room can be scheduled.', there are three checkboxes: 'Allow repeating meetings' (checked), 'Allow scheduling only during working hours' (unchecked), and 'Always decline if the end date is beyond this limit' (checked). Below these are two input fields: 'Maximum booking lead time (days):' with the value '180' and 'Maximum duration (hours):' with the value '24.0'. At the bottom, there is a text box for a reply message and two buttons: 'Save' and 'Cancel'.

FIGURE 6-4 Set restrictions for booking rooms.

The booking options are the same for both rooms and equipment. The options allow you to:

- Specify whether repeat bookings are allowed. By default, repeat bookings are allowed. If you disable the related settings, users won't be able to schedule repeating meetings.
- Specify whether the room or equipment can be scheduled only during working hours. By default, this option is disabled, which allows rooms and equipment to be scheduled for use at any time. The standard working hours are defined as 8:00 AM to 5:00 PM Monday through Friday but can be changed using the Calendaring options in Outlook.
- Specify the maximum number of days in advance the room or equipment can be booked. By default, rooms and equipment can be booked up to 180 days in advance. You can change the default to any value from 0 to 1080. A value of 0 removes the lead time restriction completely.
- Specify the maximum duration that the room or equipment can be reserved. By default, rooms and equipment can be reserved for up to 24 hours, which allows for preparation and maintenance that may be required. You can change the default to any value from 0 to 35791394.1. A value of 0 removes the duration restriction completely.

You can configure booking options after you create the room or equipment mailbox. In

Exchange Admin Center, navigate to Recipients > Resources and then double-click the resource you want to configure. Next, in the properties dialog box for the resource, select Booking Options. After you change the booking options, click Save to apply the changes.

Adding Room Mailboxes

In Exchange Admin Center, room mailboxes are displayed under Recipients > Resources. In Exchange Management Shell, you can find all room mailboxes in the organization by entering:

```
Get-Mailbox -ResultSize unlimited -Filter {(RecipientTypeDetails  
-eq 'RoomMailbox')}
```

You can create room mailboxes by completing the following steps:


1. In Exchange Admin Center, select **Recipients** in the Navigation menu and then select **Resources**.
2. Click **New** (), and then select **Room Mailbox**. This opens the New Room Mailbox dialog box, shown in Figure 6-5.
3. Type a descriptive display name in the Room Name text box.
4. For on-premises Exchange, enter the Exchange alias in the Alias text box. The Exchange alias is used to set the default email address.
5. For Exchange Online, enter the Exchange alias in the Email Address text box and then use the drop-down list to select the domain with which the room is to be associated. The Exchange Alias and the domain name are combined to set the fully qualified name, such as room4@imagedlands.onmicrosoft.com.
6. For on-premises Exchange, the Organizational Unit text box shows where in Active Directory the user account will be created. By default, this is the Users container in the current domain. Because you'll usually need to create room and equipment accounts in a specific organizational unit rather than in the Users container, click Browse to the right of the Organizational Unit text box. Use the Select Organizational Unit dialog box to choose the location in which to store the account, and then click OK.

FIGURE 6-5 Create a special mailbox for a conference room.

7. Specify the room location, phone number and capacity using the text boxes provided.
8. With on-premises Exchange, click More Options to configure these additional options:
 - **Mailbox Database** If you want to specify a mailbox database rather than use an automatically selected one, click Browse to the right of the Mailbox Database text box. In the Select Mailbox Database dialog box, choose the mailbox database in which the mailbox should be stored. Mailbox databases are listed by name as well as by associated server and Exchange version running on the server.
 - **Address Book Policy** If you've implemented address book policies to provide customized address book views, select the address book policy to associate with the equipment mailbox.
9. Click Save to create the room mailbox. If an error occurs during account or mailbox creation, neither the account nor the related mailbox will be created. You need to correct the problem before you can complete this procedure.

By default, booking requests are accepted or declined automatically based on availability. Here, the first person to reserve the room gets the reservation. If your organization has resource coordinators, you can change the booking options by completing the following steps:

1. Double-click the room mailbox with which you want to work.
2. On the Booking Delegates page, choose the Select Delegates option.
3. Next, use the options under Delegates to specify the coordinator. Click the Add

button, use the Select Delegates dialog box to select coordinators for the room. Simply double-click to add a name to the list of delegates.

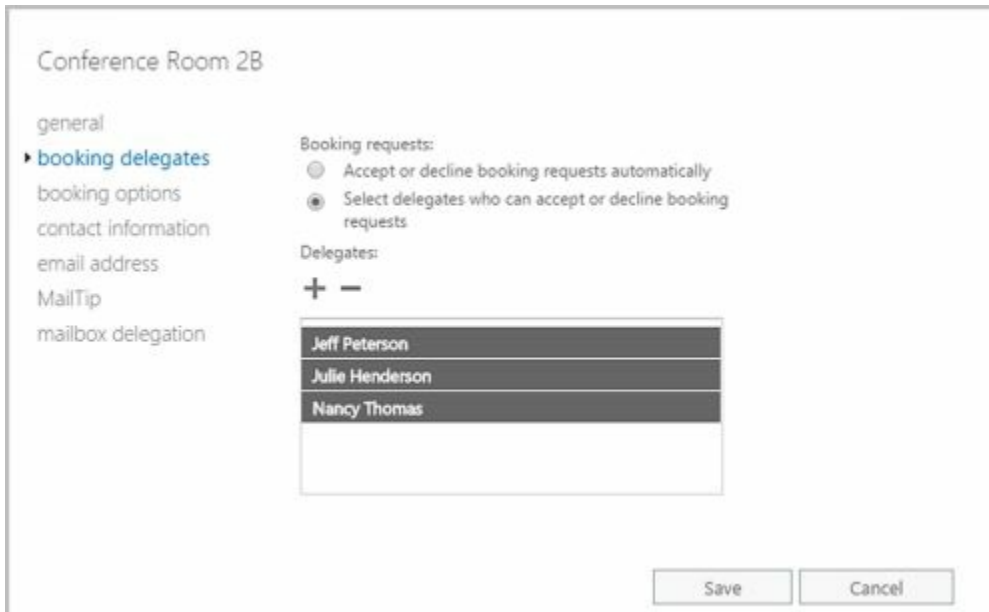


FIGURE 6-6 Add delegates if you don't want booking requests to be handled automatically.

In Exchange Management Shell, you can create a user account with a mailbox for rooms by using the `New-Mailbox` cmdlet. Sample 6-1 provides the syntax and usage.

NOTE For rooms, you must use the `-Room` parameter. For equipment, you must use the `-Equipment` parameter. By default, when you use either parameter, the related value is set as `$true`. Additionally, although with earlier releases of Exchange you needed to set a password for the related user account, this is no longer required. When you create mailboxes for Exchange Online, you cannot specify a database.

SAMPLE 6-1 Creating room mailboxes

Syntax

```
New-Mailbox -Name ' DisplayName' -Alias ' ExchangeAlias'  
-OrganizationalUnit ' OrganizationalUnit'  
-UserPrincipalName ' LogonName' -SamAccountName ' prewin2000logon'  
-FirstName ' ' -Initials ' ' -LastName ' '  
-Database ' Server \ MailboxDatabase '  
-Room
```

Usage

```
New-Mailbox -Name 'Conference Room 27' -Alias 'room27'  
-OrganizationalUnit 'imaginedlands.com/Sales'  
-UserPrincipalName 'room27@imaginedlands.com' -SamAccountName 'room27'  
-FirstName ' ' -Initials ' ' -LastName ' '  
-Database 'Sales Primary'  
-Room
```


Adding Equipment Mailboxes

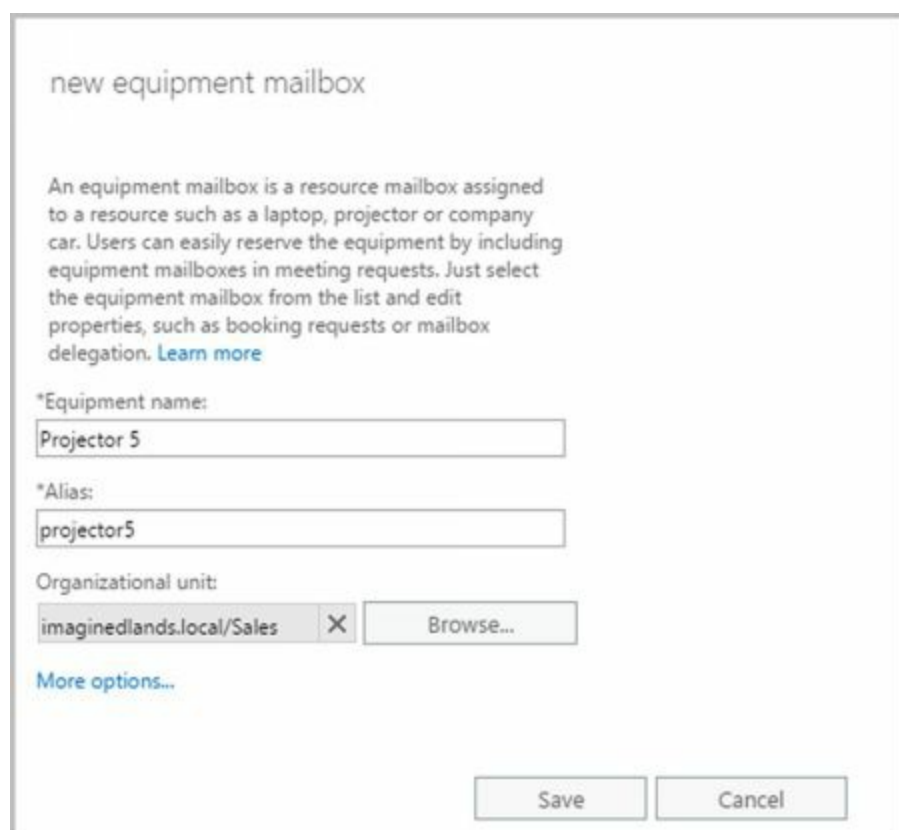
In Exchange Admin Center, equipment mailboxes are displayed under Recipients > Resources. In Exchange Management Shell, you can find all equipment mailboxes in the organization by entering:

```
Get-Mailbox -ResultSize unlimited -Filter {(RecipientTypeDetails  
-eq 'EquipmentMailbox')}
```

You can create equipment mailboxes by completing the following steps:

1. In Exchange Admin Center, select **Recipients** in the Navigation menu and then select **Resources**.

2. Click **New** (), and then select **Equipment Mailbox**. This opens the New Equipment Mailbox dialog box, shown in Figure 6-7.



new equipment mailbox

An equipment mailbox is a resource mailbox assigned to a resource such as a laptop, projector or company car. Users can easily reserve the equipment by including equipment mailboxes in meeting requests. Just select the equipment mailbox from the list and edit properties, such as booking requests or mailbox delegation. [Learn more](#)

*Equipment name:

*Alias:

Organizational unit:

[More options...](#)

FIGURE 6-7 Create a special mailbox for equipment.

3. Type a descriptive display name in the Equipment Name text box.
4. For on-premises Exchange, enter the Exchange alias in the Alias text box. The Exchange alias is used to set the default email address.
5. For Exchange Online, enter the Exchange alias in the Email Address text box and then use the drop-down list to select the domain with which the room is to be associated. The Exchange Alias and domain name are combined to set the full name, such as projector5@imaginedlands.onmicrosoft.com.
6. For on-premises Exchange, the Organizational Unit text box shows where in

Active Directory the user account will be created. By default, this is the Users container in the current domain. Because you'll usually need to create room and equipment accounts in a specific organizational unit rather than in the Users container, click **Browse** to the right of the Organizational Unit text box. Use the Select Organizational Unit dialog box to choose the location in which to store the account, and then click **OK**.

7. With on-premises Exchange, click **More Options** to configure these additional options:

- **Mailbox Database** If you want to specify a mailbox database rather than use an automatically selected one, click the Browse button to the right of the Mailbox Database text box. In the Select Mailbox Database dialog box, choose the mailbox database in which the mailbox should be stored. Mailbox databases are listed by name as well as by associated server and Exchange version running on the server.
- **Address Book Policy** If you've implemented address book policies to provide customized address book views, select the address book policy to associate with the equipment mailbox.

8. Click **Save** to create the equipment mailbox. If an error occurs during account or mailbox creation, neither the account nor the related mailbox will be created. You need to correct the problem before you can complete this procedure.

By default, booking requests are accepted or declined automatically based on availability. Here, the first person to reserve the equipment gets the reservation. If your organization has resource coordinators, you can change the booking options by completing the following steps:

1. Double-click the room mailbox with which you want to work.
2. On the Booking Delegates page, choose the Select Delegates option.
3. Next, use the options under Delegates to specify the coordinator. Click the Add button, use the Select Delegates dialog box to select coordinators for the room. Simply double-click to add a name to the list of delegates.

In Exchange Management Shell, you can create a user account with a mailbox for equipment by using the New-Mailbox cmdlet. Sample 6-2 provides the syntax and usage. Although with earlier releases of Exchange you needed to set a password for the related user account, this is no longer required. When you create mailboxes for Exchange Online, you cannot specify a database.

SAMPLE 6-2 Creating equipment mailboxes

Syntax

```
New-Mailbox -Name ' DisplayName' -Alias ' ExchangeAlias'  
-OrganizationalUnit ' OrganizationalUnit'  
-UserPrincipalName ' LogonName' -SamAccountName ' prewin2000logon'  
-FirstName ' ' -Initials ' ' -LastName ' '  
-Database ' Server \ MailboxDatabase '
```

-Equipment

Usage

New-Mailbox -Name 'Media Cart 3' -Alias 'cart3'

-OrganizationalUnit 'imaginedlands.com/Marketing'

-UserPrincipalName 'cart3@imaginedlands.com' -SamAccountName 'cart3'

-FirstName " -Initials " -LastName "

-Database 'Marketing Primary'

-Equipment

Adding Linked Mailboxes

A linked mailbox is a mailbox that is accessed by a user in a separate, trusted forest. Typically, you use linked mailboxes when your organization's mailbox servers are in a separate resource forest and you want to ensure that users can access free/busy data across these forests. You use linked mailboxes with on-premises Exchange organizations.


All linked mailboxes have two user account associations:

- A unique user account in the same forest as the Mailbox server. The same forest user account is disabled automatically so that it cannot be used for logon.
- A unique user account in a separate forest for which you are creating a link. The separate forest user account is enabled so that it can be used for logon.

In Exchange Admin Center, linked mailboxes are displayed under Recipients > Mailboxes. In Exchange Management Shell, you can find all linked mailboxes in the organization by entering:

```
Get-Mailbox -ResultSize unlimited -Filter {(RecipientTypeDetails -eq 'LinkedMailbox')}
```

You can create a linked mailbox by completing the following steps:

1. In Exchange Admin Center, select **Recipients** in the Navigation menu, and then select **Mailboxes**.
2. Click **New** (), and then select **Linked Mailbox**. This starts the New Linked Mailbox Wizard. A linked mailbox cannot be created without a forest or domain trust in place between the source and target forests.
3. On the New Linked Mailbox page, click **Browse** to the right of the Linked Forest text box. In the Select Trusted Forest Or Domain dialog box, select the linked forest or domain in which the user's original account is located, and then click **OK**. This is the separate forest that contains the user account that you want to create the linked mailbox for in the current forest. Click **Next**.
4. If your organization has configured a one-way outgoing trust where the current forest trusts the linked forest, you're prompted for administrator credentials in the linked forest so that you can gain access to a domain controller in that forest. Type the user name and password for an administrator account in the account forest, and then click **Next**.
5. Click **Browse** to the right of the Linked Domain Controller text box. In the Select Domain Controller dialog box, select a domain controller in the linked forest, and then click **OK**.
6. Click **Browse** to the right of the Linked Master Account text box. Use the options in the Select User dialog box to select the original user account in the linked forest, and then click **OK**.

7. Click **Next**. On the General Information page, the Organizational Unit text box shows where in Active Directory the user account will be created. By default, this is the Users container in the current domain. Select the Specify The Organizational Unit check box and then click Browse to create the new user account in a different container. Use the Select Organizational Unit dialog box to choose the location in which to store the account, and then click OK.
8. In the User Logon Name text box, type the user's logon name. Use the drop-down list to select the domain with which the account is to be associated. This sets the fully qualified logon name.
9. Click **More Options**. Type the user's first name, middle initial, and last name in the text boxes provided. These values are used to create the Name entry, which is the user's display name.
10. Optionally, enter an Exchange alias for the user. The alias must be unique in the forest. If you don't specify an alias, the logon name is used as the alias.
11. If you want to specify a mailbox database rather than use an automatically selected one, click Browse to the right of the Mailbox Database text box. In the Select Mailbox Database dialog box, choose the mailbox database in which the mailbox should be stored. Mailbox databases are listed by name as well as by associated server and Exchange version running on the server.
12. Click **Save** to create the account and the related mailbox. If an error occurs during account or mailbox creation, neither the account nor the related mailbox will be created. You will need to correct the problem.

In Exchange Management Shell, you can create a user account with a linked mailbox by using the New-Mailbox cmdlet. Sample 6-3 provides the syntax and usage. You'll be prompted for the credentials of an administrator account in the linked forest. Although with earlier releases of Exchange you needed to set a password for the related user account, this is no longer required.

SAMPLE 6-3 Creating linked mailboxes

Syntax

```
New-Mailbox -Name ' DisplayName' -Alias ' ExchangeAlias'
-OrganizationalUnit ' OrganizationalUnit'
-Database ' Database'
-UserPrincipalName ' LogonName' -SamAccountName ' prewin2000logon'
-FirstName ' FirstName' -Initials ' Initial' -LastName ' LastName'
-ResetPasswordOnNextLogon State
-LinkedDomainController ' LinkedDC'
-LinkedMasterAccount ' domain\user'
-LinkedCredential:(Get-Credential ' domain\administrator' )
```

Usage

```
New-Mailbox -Name 'Wendy Richardson' -Alias 'wendyr'
```



```
-OrganizationalUnit 'imaginedlands.com/Sales'  
-Database 'Corporate Services Primary'  
-UserPrincipalName 'wendyr@imaginedlands.com' -SamAccountName 'wendyr'  
-FirstName 'Wendy' -Initials '' -LastName 'Richardson'  
-ResetPasswordOnNextLogon $true  
-LinkedDomainController 'TvpresDC58'  
-LinkedMasterAccount 'tvpres\wrichardson'  
-LinkedCredential:(Get-Credential 'tvpres\williams' )
```

Working with Archive Mailboxes

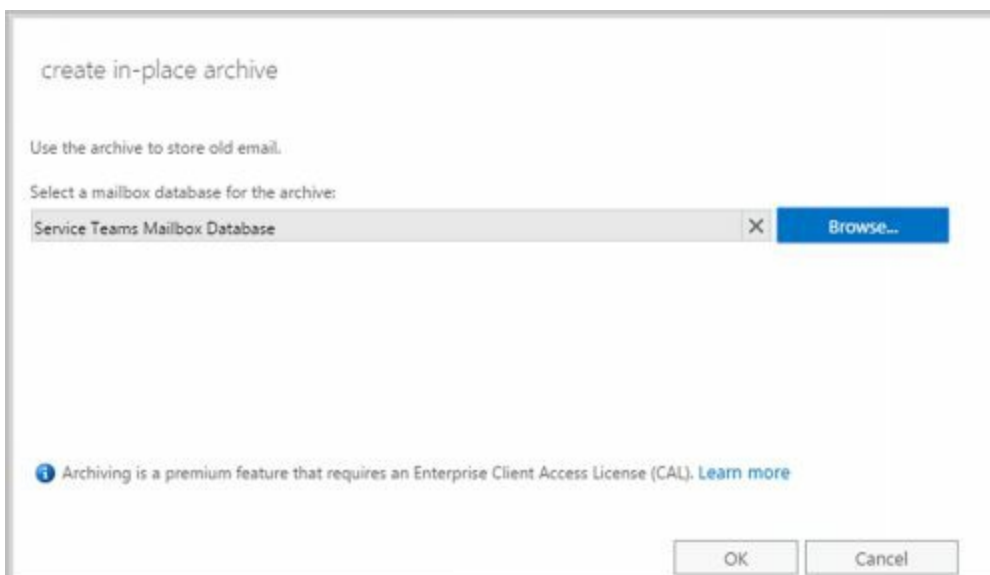
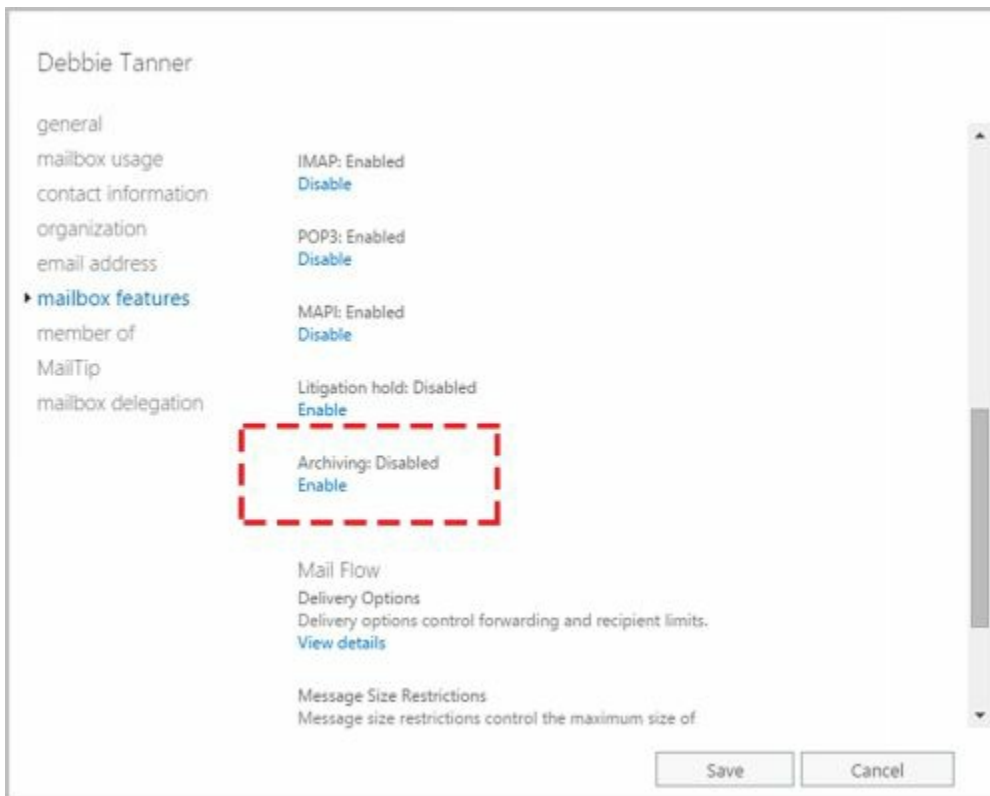
Each user can have an alternate mailbox for archives. An archive mailbox is used to store a user's old messages, such as might be required for executives and needed by some managers and users. In Outlook and Outlook Web App, users can access archive mailboxes in much the same way as they access a regular mailbox.

Archive mailboxes are created in one of two ways. The standard approach is to create an in-place archive. Both on-premises Exchange and Exchange Online use in-place archives by default. With hybrid organizations, you also can use online archives. With an online archive, the archive for an on-premises mailbox is created in the online service.

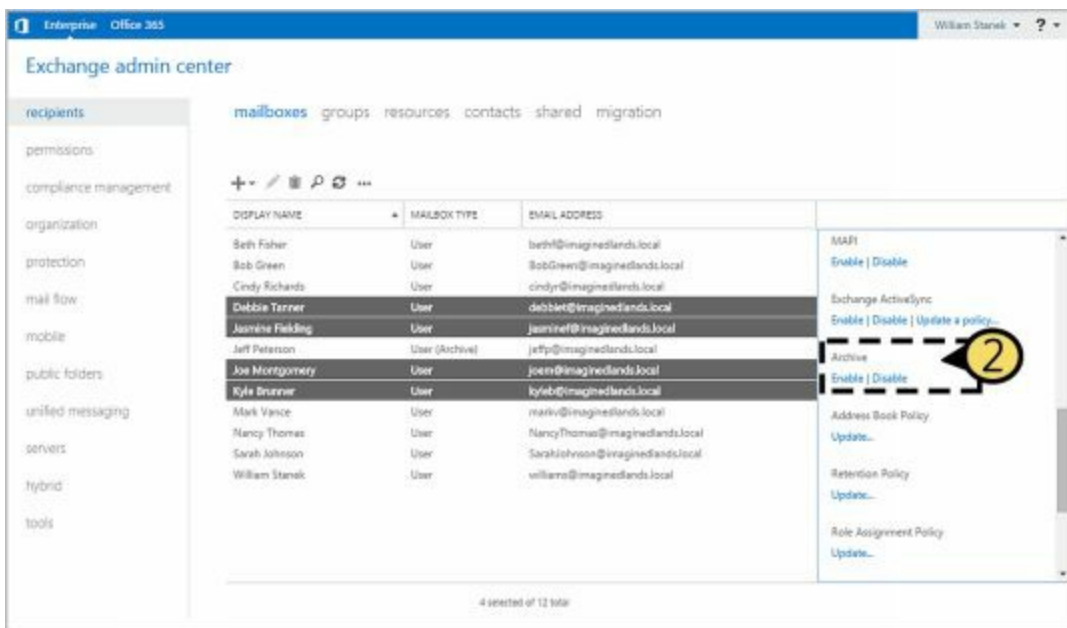
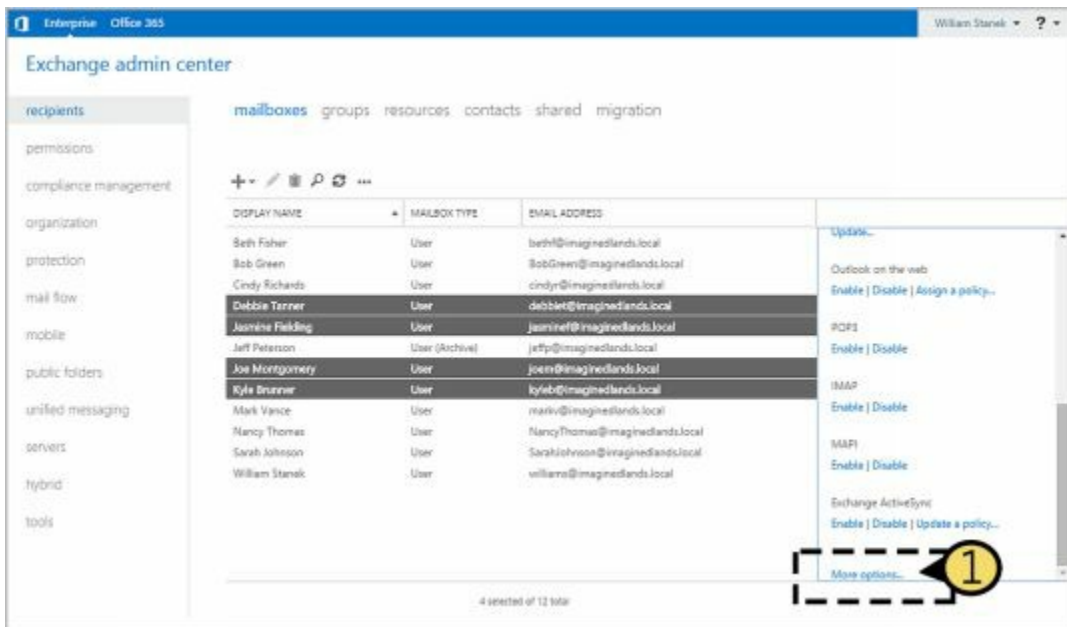
Adding In-Place Archives

You can create an in-place archive mailbox at the same time you create the user's standard mailbox. To create an in-place archive mailbox, complete the following steps:

1. In Exchange Admin Center, select **Recipients** in the Navigation menu and then select **Mailboxes**. Double-click the entry for the user's standard mailbox. Any user that already has an archive mailbox has "User (Archive)" as the mailbox type.
2. On the Mailbox Features page, the status of archiving is listed under the Archiving heading. If archiving is disabled, select **Enable** under the Archiving heading and continue with this procedure.
3. With on-premises Exchange, if the mailbox had an archive previously and that archive still exists, this archive is used in its original location. Otherwise, the Create In-Place Archive dialog box is displayed. If you want to specify a mailbox database rather than use an automatically selected one, click **Browse** to the right of the Mailbox Database text box. In the Select Mailbox Database dialog box, choose the mailbox database in which the mailbox should be stored and then click **OK**. Mailbox databases are listed by name as well as by associated server and Exchange version running on the server.
4. Click **Save**. If an error occurs during mailbox creation, the archive mailbox will not be created. You need to correct the problem before you can complete this procedure and create the archive mailbox.



When you are working with Exchange Admin Center, you can enable in-place archiving for multiple mailboxes as well. When you select multiple mailboxes using the Shift or Ctrl keys, the Details pane displays bulk editing options. Scroll down the list of options and then click **More Options**. Next, under Archive, click **Enable**.



The Bulk Enable Archive dialog box is displayed. If you want to specify a mailbox database for the archives rather than use an automatically selected one, click **Browse** to the right of the Mailbox Database text box. In the Select Mailbox Database dialog box, choose the mailbox database in which the archive mailboxes should be stored and then click **OK**.

Using Exchange Management Shell, you can create an archive mailbox using the Enable-Mailbox cmdlet. The basic syntax is as follows:

Enable-Mailbox [-Identity] **Identity** –Archive [-Database **DatabaseID**]

such as:

Enable-Mailbox imaginedlands.com/engineering/tonyg –archive

Because each user can have only one archive mailbox, you get an error if the user already has an archive mailbox. Items in the user’s mailbox will be moved automatically to the archive mailbox based on the archive and retention policy. When you install Exchange Server, a default archive and retention policy is created for all

archive mailboxes. This policy is named Default MRM Policy. Because of this policy, email messages from the entire mailbox are moved to the archive after two years by default.

For bulk editing, you can use various techniques. Generally, you'll want to:

- Ensure you are working with mailboxes for regular users and not mailboxes for rooms, equipment, and so on. To do this, filter the results based on the `RecipientTypeDetails`.
- Ensure the mailbox doesn't already have an on-premises or online archive. To do this, filter the results based on whether the mailbox has an associated `ArchiveGuid` and the `ArchiveDomain`.
- Ensure you don't enable archives on mailboxes that shouldn't have them, such as the Discovery Search Mailbox. To do this, filter based on the name or partial name of mailboxes to exclude.

Consider the following example:

```
Get-Mailbox -Database Sales -Filter {RecipientTypeDetails -eq 'UserMailbox'  
-AND ArchiveGuid -eq $null -AND ArchiveDomain -eq $null -AND Name -NotLike  
"DiscoverySearchMailbox*"} | Enable-Mailbox -Archive
```

In this example, `Get-Mailbox` retrieves all mailboxes for regular users in the Sales database that don't have in-place or online archiving enabled and that also don't have a name starting with: `DiscoverySearchMailbox`. The results are then piped through `Enable-Mailbox` to add an archive mailbox to these mailboxes.

Adding Online Archives

In hybrid organizations, several features, including online archives, are enabled by default. If you are unsure whether online archives have been enabled for your hybrid deployment, enter **Get-HybridConfiguration | fl** at a PowerShell prompt and then verify that the `OnlineArchive` flag is set on the `Features` parameter. To modify the hybrid configuration, you can use `Set-HybridConfiguration`. However, do not use `Set-HybridConfiguration` without a solid understanding of hybrid configurations. Keep in mind that when you use the `-Features` parameter with `Set-HybridConfiguration`, you must explicitly specify all the features that you want enabled. Any feature that you omit will be disabled.

In Exchange Management Shell, you create online archives using the `Enable-Mailbox` cmdlet with the `-RemoteArchive`, `-ArchiveDatabase`, and `-ArchiveDomain` parameters. The required `-RemoteArchive` parameter is a flag that specifies you want to create the archive online. The optional `-ArchiveDatabase` sets the name or GUID of the archive database in the online organization. The optional `-ArchiveDomain` sets the fully qualified domain name of the domain for the online organization. Consider the following examples:

```
Enable-Mailbox -Identity issan@contoso.com -RemoteArchive
```

```
Enable-Mailbox -Identity issan@contoso.com -RemoteArchive -ArchiveDatabase  
"D919BA05-46A6-415f-80AD-7E09334BB852" -ArchiveDomain  
"imaginedlands.onmicrosoft.com"
```

The first example creates the online archive using the default database and online domain. The second example explicitly sets the GUID of the database and domain parameters.

Managing Archive Settings

Whether you use Exchange Admin Center or Exchange Management Shell, several other parameters are set for archive mailboxes. The default name for the archive mailbox is set as In-Place Archive – *UserDisplayName*, such as In-Place Archive – Henrik Larsen. With on-premises Exchange, the default quota and warning quota are set as 50 GB and 45 GB respectively. With Exchange Online, the default quota and warning quota are set as 25 GB and 22.5 GB, respectively.

You can confirm the details for a user's archive mailbox by entering the following command:

```
Get-Mailbox "Name" | fl name, alias, servername, *archive*
```

where *name* is the display name or alias of the user you want to work with, such as:

```
Get-Mailbox "Henrik Larsen" | fl name, alias, servername, *archive*
```

You can change the archive name and set quotas by using Set-Mailbox. The basic syntax is as follows:

```
Set-Mailbox[-Identity] Identity –ArchiveName Name  
-ArchiveQuota Quota -ArchiveWarningQuota Quota
```

When you set a quota, specify the value with MB (for megabytes), GB (for gigabytes), or TB (for terabytes), or enter 'Unlimited' to remove the quota. Here is an example:

```
set-mailbox imaginedlands.com/engineering/tonyg  
-ArchiveQuota '28GB' -ArchiveWarningQuota '27GB'
```

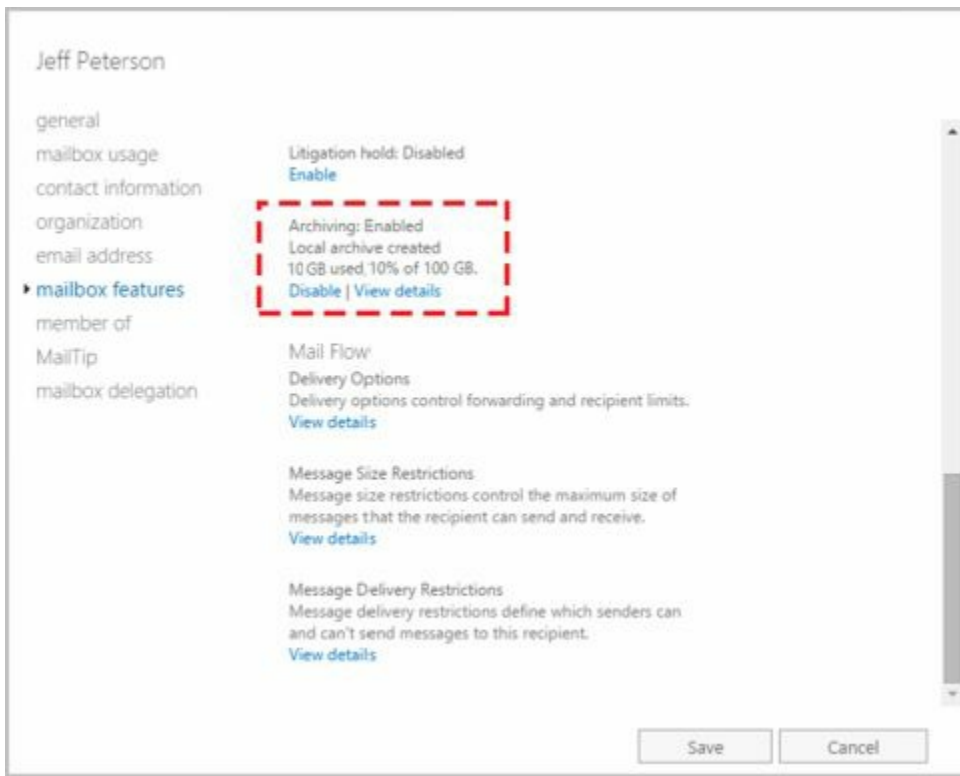
For bulk editing, you can use Get-Mailbox to retrieve the user mailboxes you want to work with and then apply the changes by piping the results to Set-Mailbox. If you do so, ensure that you filter the results appropriately. Consider the following example:

```
Get-Mailbox -ResultSize unlimited -Filter {RecipientTypeDetails -eq  
'UserMailbox' -AND ArchiveGuid -ne $null} | Set-Mailbox -ArchiveQuota  
'20GB' -ArchiveWarningQuota '18GB'
```

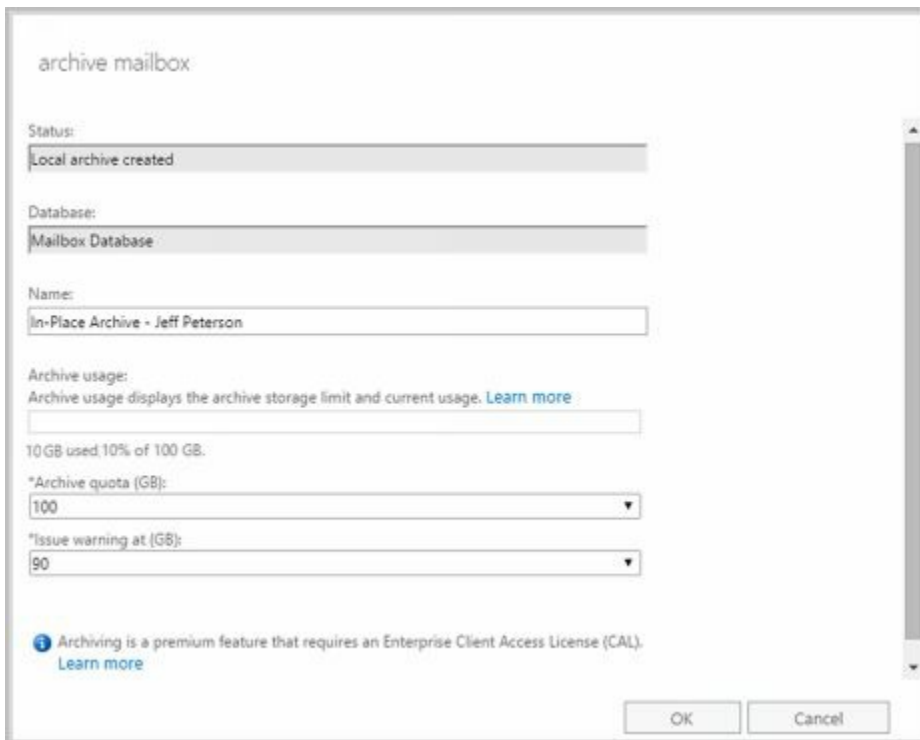
In this example, Get-Mailbox retrieves all mailboxes for regular users in the entire organization that have archiving enabled. The results are then piped through Set-Mailbox to modify the quota and quota warning values.

In Exchange Admin Center, you manage archive settings by completing these steps:

1. In Exchange Admin Center, select Recipients in the Navigation menu and then select Mailboxes. Double-click the entry for the user's standard mailbox. Any user that already has an archive mailbox has "User (Archive)" as the mailbox type.
2. On the Mailbox Features page, click **View Details** under the Archiving heading.

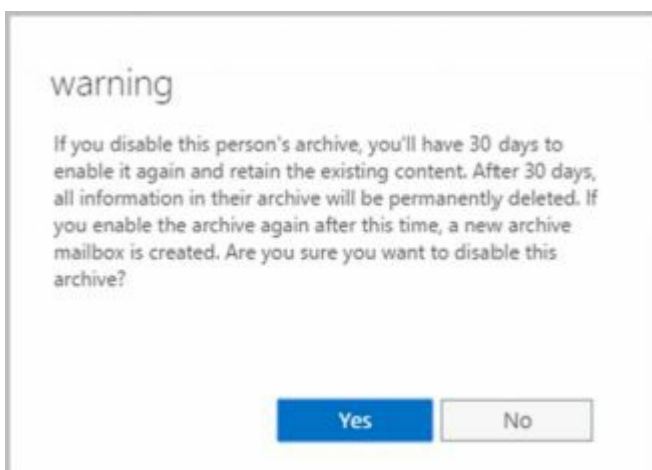


3. To change the name of the archive mailbox, enter the new name in the Name text box.
4. To set a quota, enter the desired value in gigabytes in the Archive Quota combo box.
5. To set a quota warning, enter a quota warning in gigabytes in the Issue Warning At combo box.

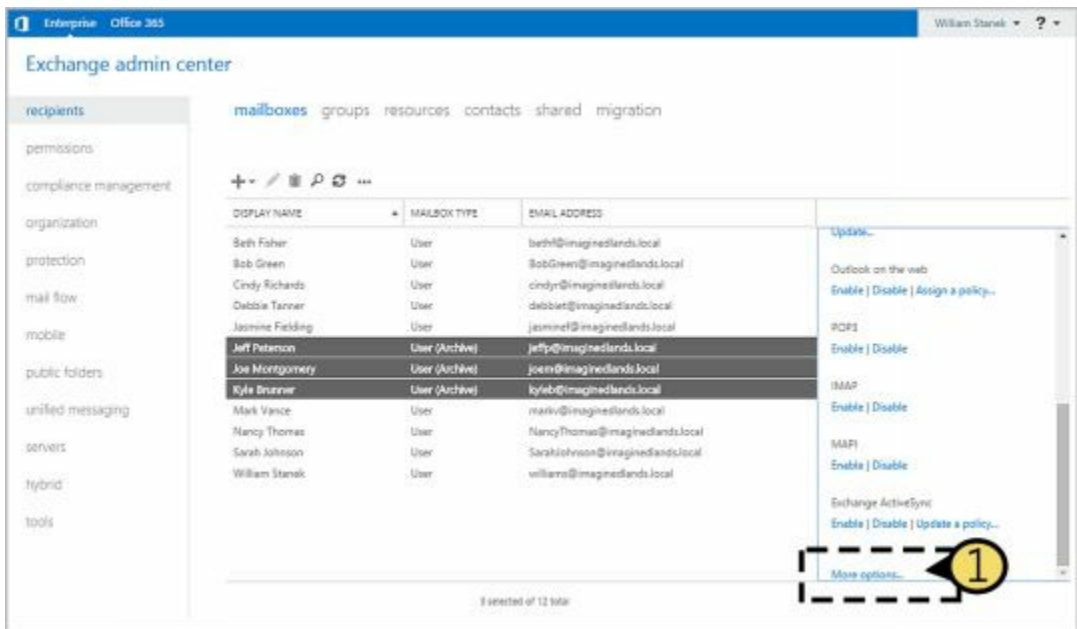


REAL WORLD When you disable an archive mailbox for a user, the archive mailbox is marked for deletion and disconnected from the user account. The archive mailbox is retained according to the mailbox retention policy. To connect the disabled archive mailbox to the existing mailbox, you must use the Connect-Mailbox cmdlet with the -Archive parameter. Otherwise, if you disable an archive mailbox for a user mailbox and then enable an archive mailbox for that same user, a new archive mailbox is created for the user.

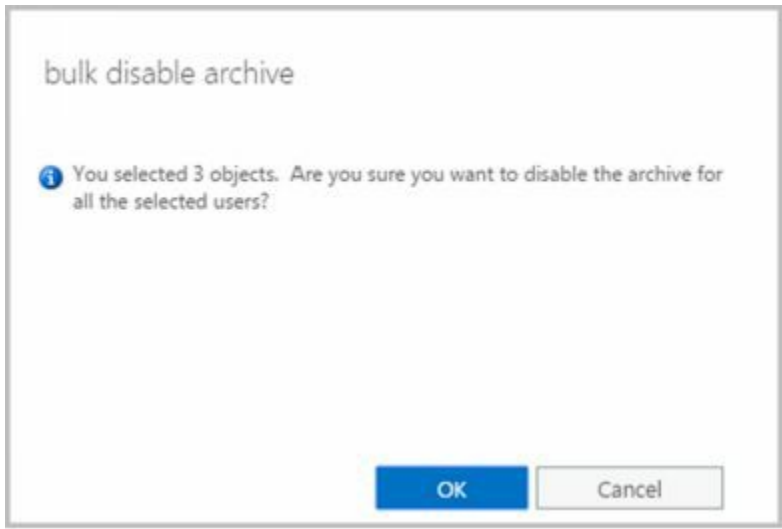
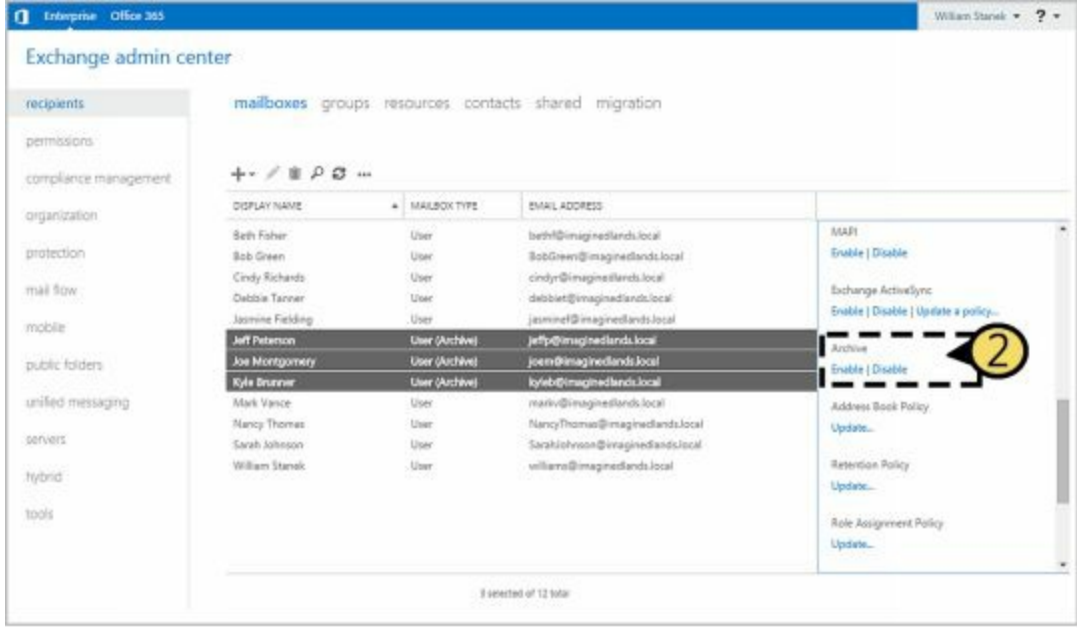
To disable an archive mailbox, open the properties dialog box for the user to the Mailbox Features page and then select **Disable** under the Archiving heading. Click **Yes** when prompted to confirm.



When you are working with Exchange Admin Center, you can disable in-place archiving for multiple mailboxes as well. When you select multiple mailboxes using the Shift or Ctrl keys, the Details pane displays bulk editing options. Scroll down the list of available options and then click **More Options**.



Next, under Archive, click **Disable**. When the Bulk Disable Archive dialog box is displayed, click **OK**.



In Exchange Management Shell, you can disable an archive mailbox by using Disable-Mailbox. The basic syntax is as follows:

Disable-Mailbox [-Identity] **Identity** –Archive

such as:

```
disable-mailbox imaginedlands.com/engineering/tonyg –archive
```

For bulk editing, you can use a technique similar to the one discussed for enabling archives. Consider the following example:

```
Get-Mailbox -Database Sales -Filter {RecipientTypeDetails -eq 'UserMailbox'  
-AND ArchiveGuid -ne $null} | Disable-Mailbox -Archive
```

In this example, Get-Mailbox retrieves all mailboxes for regular users in the Sales database that have archiving enabled. The results are then piped through Disable-Mailbox to remove the archive mailbox from these mailboxes.

Adding Arbitration Mailboxes

Exchange moderated transport requires all email messages sent to specific recipients to be approved by moderators. You can configure any type of recipient as a moderated recipient, and Exchange will ensure that all messages sent to those recipients go through an approval process.

Distribution groups are the only types of recipients that use moderation by default. Membership in distribution groups can be closed, owner approved, or open. While any Exchange recipient can join or leave an open distribution group, joining or leaving a closed group requires approval. Group owners receive join and remove requests and can either approve or deny those requests.

Distribution groups can also be unmoderated or moderated. With unmoderated groups, any approved sender (which is all senders by default) can send messages to the group. With moderated groups, messages are sent to moderators for approval before being distributed to members of the group. The only exception is for a message sent by a moderator. A message from a moderator is delivered immediately because a moderator has the authority to determine what is and isn't an appropriate message.

NOTE The default moderator for a distribution group is the group's owner.

Arbitration mailboxes are used to store messages that are awaiting approval. When you install Exchange Server 2016, a default arbitration mailbox is created. For the purposes of load balancing or for other reasons, you can convert other mailboxes to the arbitration mailbox type by using the Enable-Mailbox cmdlet. The basic syntax is as follows:

```
Enable-Mailbox [-Identity] Identity -Arbitration
```

such as:

```
enable-mailbox imagedlands.com/users/moderatedmail -Arbitration
```

You can create an arbitration mailbox by using New-Mailbox as shown in this example:

```
New-Mailbox ModeratedMail -Arbitration -UserPrincipalName  
ModeratedMail@imagedlands.com
```

Adding Discovery Mailboxes

Exchange Discovery helps organizations comply with legal discovery requirements and can also be used as an aid in internal investigations or as part of regular monitoring of email content. Exchange Discovery uses content indexes created by Exchange Search to speed up the search process.

NOTE By default, Exchange administrators do not have sufficient rights to perform Discovery searches. Only users with the Discovery Management role can perform Discovery searches. If a user is not a member of the role, she doesn't have access to the related options. This means she can't access the In-Place eDiscovery & Hold interface in Exchange Admin Center or the In-Place eDiscovery & Hold cmdlets in PowerShell.

Whether you are working in an online, on-premises, or hybrid organization, you use Exchange Admin Center to perform searches. With hybrid configurations, an on-premises search will return results from the online organization.

Discovery searches are performed against designated mailboxes or all mailboxes in the Exchange organization. Items in mailboxes that match the Discovery search are copied to a target mailbox. Only mailboxes specifically designated as Discovery mailboxes can be used as targets. In a hybrid configuration, you must copy items to an on-premises mailbox, regardless of whether the items are from the online or on-premises organization.

TIP By default, Discovery search does not include items that cannot be indexed by Exchange Search. To include such items in the search results, select the Include Items That Can't Be Searched check box in Exchange Admin Center.

In Exchange Admin Center, you can access the discovery and hold settings by selecting Compliance Management in the Navigation menu and then selecting In-Place eDiscovery & Hold. While working with In-Place eDiscovery & Hold, you can create searches across mailboxes by specifying filters and hold options for search results.

When you install Exchange Server 2016, a default Discovery mailbox is created. You can convert other mailboxes to the Discovery mailbox type by using the Enable-Mailbox cmdlet. The basic syntax is as follows:

```
Enable-Mailbox [-Identity] Identity -Discovery
```

such as:

```
enable-mailbox imaginedlands.com/hr/legalsearch -discovery
```

You can create a Discovery mailbox by using New-Mailbox as shown in this example:

```
New-Mailbox LegalSearch -Discovery -UserPrincipalName  
LegalSearch@imaginedlands.com
```

Once a Discovery mailbox is established, you can't convert it to another mailbox type. You can't use Exchange Admin Center to create Discovery mailboxes.

Adding Shared Mailboxes

Shared mailboxes are mailboxes that are shared by multiple users. Although shared mailboxes must have an associated user account, this account is not used for logon in the domain and is disabled by default. Users who access the shared mailbox do so using access permissions.

You can create a shared mailbox by using New-Mailbox, as shown in this example:

```
New-Mailbox -Shared -Name "Customer Service" -DisplayName  
"Customer Service" -Alias Service -UserPrincipalName  
customerservice@imaginedlands.com
```

In this example, a user account named CustomerService is created for this mailbox. This user account is disabled by default to prevent logon using this account. After creating the mailbox, you need to grant Send On Behalf Of permission to the appropriate users or security groups by using Set-Mailbox and the -GrantSendOnBehalfTo parameter. Finally, you need to add access rights that allow these users or security groups to log on to the mailbox by using Add-MailboxPermission and the -AccessRights parameter. Ensure these rights are inherited at all levels of the mailbox using -InheritanceType All as well. One way this would all come together is shown in the following example:

```
New-Mailbox -Shared -Name "Customer Service" -DisplayName  
"Customer Service" -Alias Service -UserPrincipalName  
customerservice@imaginedlands.com | Set-Mailbox -GrantSendOnBehalfTo  
CustomerServiceGroup | Add-MailboxPermission -User CustomerServiceGroup  
-AccessRights FullAccess -InheritanceType All
```

In Exchange Admin Center, you can create a shared mailbox by following these steps:

1. Select **Recipients** in the Navigation menu and then select **Shared**.

2. Click **New** (). This opens the New Shared Mailbox dialog box, shown in Figure 6-8.

new shared mailbox

Shared mailboxes allow a group of users to view and send email from a common mailbox and share a common calendar. [Learn more](#)

*Display name:
Customer Service

*Alias:
service

Organizational unit:
imaginedlands.local/Custon X Browse...

Users
The following users have permission to view and send mail from this shared mailbox.

+ -

DISPLAY NAME ▲


Cindy Richards

Jasmine Fielding

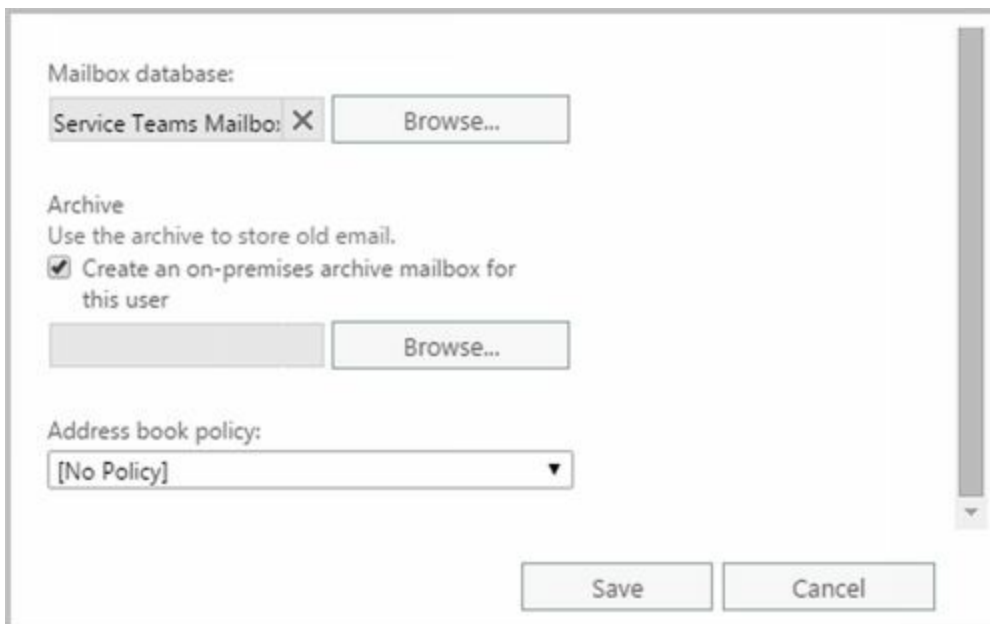
[More options...](#)

Save Cancel

FIGURE 6-8 Create a mailbox to share with multiple users.

3. In the Display Name text box, type a descriptive name for the shared mailbox.
4. For on-premises Exchange, the Organizational Unit text box shows where in Active Directory the associated user account will be created. By default, this is the Users container in the current domain. If you want to use a different container, click Browse to the right of the Organizational Unit text box. Use the Select Organizational Unit dialog box to choose the location in which to store the account, and then click OK.
5. For on-premises Exchange, enter the Exchange alias in the Alias text box. The Exchange alias is used to set the default email address.
6. For Exchange Online, enter the Exchange alias in the Email Address text box and then use the drop-down list to select the domain with which the room is to be associated. The Exchange Alias and domain name are combined to set the fully qualified name, such as `service@imaginedlands.onmicrosoft.com`.
7. Under Users, click Add (). In the Select Full Access dialog box, select users, security groups, or both that should be able to view and send email from the shared mailbox. Select multiple users and groups using the Shift or Ctrl keys.
8. Click **More Options** to configure these additional options:

- **Alias** For Exchange Online, sets the Exchange alias and overrides the default value you set previously using the Email Address text box. This allows a resource to have an alias that is different from the name portion of its email address.
 - **Mailbox Database** For on-premises Exchange, if you want to specify a mailbox database rather than use an automatically selected one, click Browse to the right of the Mailbox Database text box. In the Select Mailbox Database dialog box, choose the mailbox database in which the mailbox should be stored. Mailbox databases are listed by name as well as by associated server and Exchange version running on the server.
 - **Archive** For on-premises Exchange, if you want to create an on-premises archive mailbox as well, select the related checkbox. Optionally, click Browse to choose the mailbox database for the archive.
 - **Address Book Policy** For on-premises Exchange, if you've implemented address book policies to provide customized address book views, select the address book policy to associate with the equipment mailbox.
9. Click **Save** to create the shared mailbox. If an error occurs during account or mailbox creation, neither the account nor the related mailbox will be created. You need to correct the problem before you can complete this procedure.



The screenshot shows a dialog box for configuring mailbox settings. It has three main sections: 'Mailbox database', 'Archive', and 'Address book policy'. At the bottom are 'Save' and 'Cancel' buttons.

Mailbox database:
Service Teams Mailbo: X Browse...

Archive
Use the archive to store old email.
 Create an on-premises archive mailbox for this user
Browse...

Address book policy:
[No Policy] ▼

Save Cancel

Adding Public Folder Mailboxes

Public folders are used to share messages and files in an organization. Public folder trees define the structure of an organization's public folders. You can make the default public folder tree accessible to users based on criteria you set, and then users can create folders and manage their content.

Each public folder in the default public folder tree can have specific access permissions. For example, you can create public folders called CompanyWide, Marketing, and Engineering. Whereas you would typically make the CompanyWide folder accessible to all users, you would make the Marketing folder accessible only to users in the marketing department and the Engineering folder accessible only to users in the engineering department.

Users access public folders from Outlook clients, including Outlook Web App and Outlook 2010 or later. With Outlook Web App and Outlook 2010 or later, users can add and remove favorite public folders and perform item-level operations, such as creating and managing posts. However, users can create or delete public folders only from Outlook 2010 or later. As an administrator, you can manage public folders in Exchange Admin Center.

Unlike Exchange 2010 and earlier versions of Exchange, current Exchange servers no longer use public folder databases or store public folder data separately from mailbox data. Instead, Exchange Server and Exchange Online store public folder data in mailboxes. This significant architecture change greatly simplifies public folder management.

In Exchange Admin Center, you work with public folders by selecting Public Folders in the Navigation menu and then selecting either Public Folder Mailboxes or Public Folders as appropriate. You use the options under Public Folder Mailboxes to create and manage the mailboxes that store public folder data. You use the options under Public Folders to view and manage the public folder hierarchy.

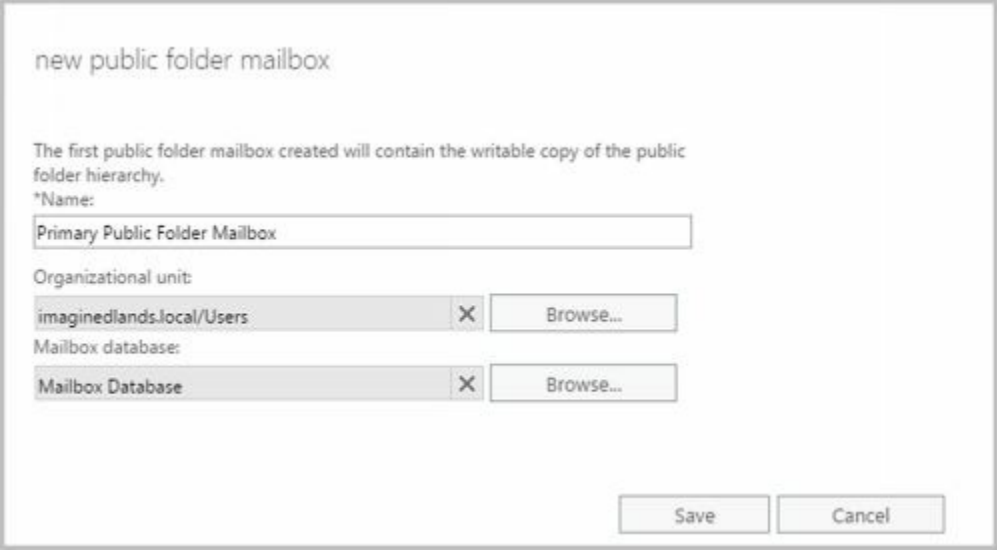
An Exchange organization can have one or more public folder mailboxes and those mailboxes can be created on one or more Mailbox servers throughout the organization. While each public folder mailbox can contain public folder content, only the first public folder mailbox created in an Exchange organization contains the writable copy of the public folder hierarchy. This mailbox is referred to as the hierarchy mailbox. Any additional public folder mailboxes contain read-only copies of the public-folder hierarchy.

Because there's only one writeable copy of the public folder hierarchy, proxying is used to relay folder changes to the hierarchy mailbox. This means that any time users working with folders in an additional mailbox create new subfolders, the folder creation, modification, or removal is proxied to the hierarchy mailbox by the content mailbox users are connected to.

In Exchange Admin Center, you can create a public folder mailbox by following these steps:

1. Select **Public Folders** in the Navigation menu and then select **Public Folder Mailboxes**.

2. Click **New** (). This opens the New Public Folder Mailbox dialog box, shown in Figure 6-9.



new public folder mailbox

The first public folder mailbox created will contain the writable copy of the public folder hierarchy.

*Name:
Primary Public Folder Mailbox

Organizational unit:
imaginedlands.local/Users X Browse...

Mailbox database:
Mailbox Database X Browse...

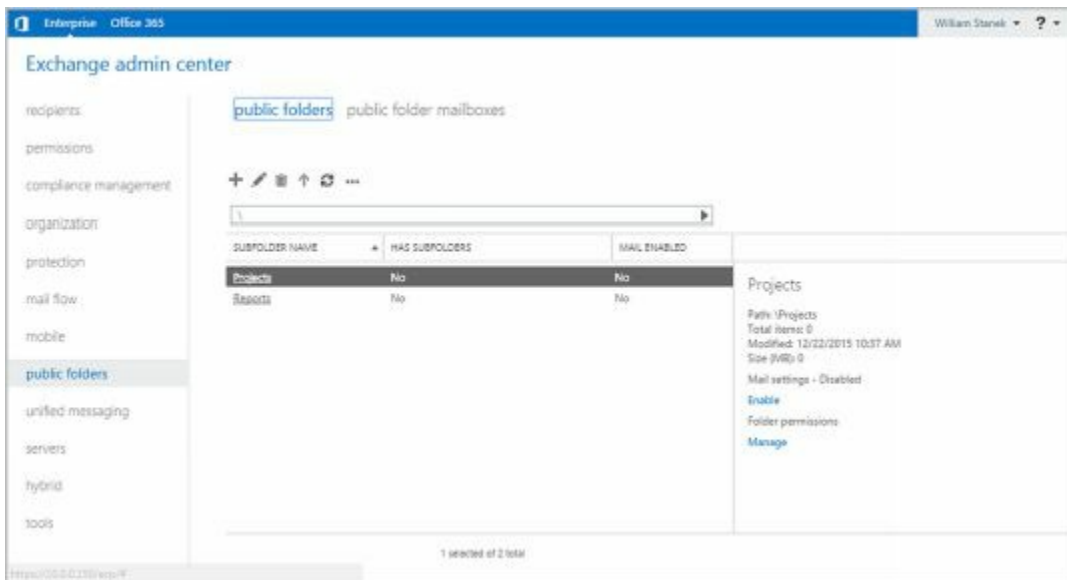
Save Cancel

FIGURE 6-9 Create a mailbox for public folder storage.

3. Type a descriptive name for the mailbox.
4. With on-premises Exchange, you can associate the mailbox with a specific organizational unit. Click **Browse** to the right of the Organizational Unit text box. Use the **Select Organizational Unit** dialog box to choose the location in which to store the account, and then click **OK**. A user account for the mailbox is created in the selected organizational unit (with the account disabled for login).
5. With on-premises Exchange, you can specify a mailbox database rather than use an automatically selected one, click **Browse** to the right of the Mailbox Database text box. In the **Select Mailbox Database** dialog box, choose the mailbox database in which the mailbox should be stored and then click **OK**.
6. Click **Save** to create the public folder mailbox. If an error occurs during account or mailbox creation, neither the account nor the related mailbox will be created. You need to correct the problem before you can complete this procedure.

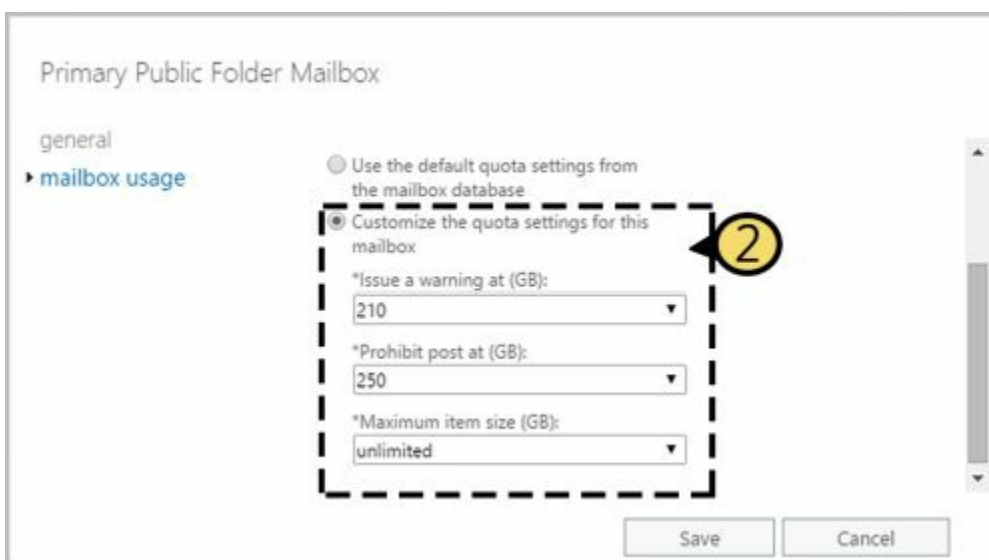
Public folder content can include email messages, documents, and more. The content is stored in the public folder mailbox but isn't replicated across multiple public folder mailboxes. Instead, all users access the same public folder mailbox for the same set of content.

When you create the first public folder in the organization, you establish the root of the public folder hierarchy. You can then create subfolders and assign access permissions on folders. In Exchange Admin Center, select **Public Folders > Public Folders** and then use the available options to create subfolders and set permissions on those folders.



When you create public folder mailboxes, they inherit the quota limits of the mailbox database in which they are stored. You can modify the quota limits using the properties dialog for the mailbox.

Double-click the mailbox entry. In the Public Folder Mailbox dialog box, on the Mailbox Usage page, click **More Options** and then select **Customize The Settings For This Mailbox**. Next, use the selection lists provided to specify when warnings are issued, what posts are prohibited, and the maximum size of items. Apply the changes by clicking **Save**.



When users are connected to public folder mailboxes and make routine changes to an Exchange store hierarchy or content, the changes are synchronized every 15 minutes using Incremental Change Synchronization (ICS). Immediate syncing is used for non-routine changes, such as folder creation. If no users are connected to public folder mailboxes, synchronization occurs once every 24 hours by default.

Chapter 7. Managing Mailboxes

The difference between a good Microsoft Exchange Server administrator and a great one is the attention he or she pays to mailbox administration. Mailboxes are private storage places for messages you've sent and received, and they are created as part of private mailbox databases in Exchange. Mailbox settings control mail delivery, permissions, and storage limits.

You can configure most mailbox settings on a per-mailbox basis. However, with Exchange Online, some settings are configured for all users of the service while other settings are fixed as part of the service and cannot be changed. With on-premises Exchange, you cannot change some settings without moving mailboxes to another mailbox database or changing the settings of the mailbox database itself. For example, with on-premises Exchange, you set the storage location on the file system, the storage limits, the deleted item retention, and the default offline address book on a per-mailbox-database basis. Keep this in mind when performing capacity planning and when deciding which mailbox location to use for a particular mailbox.

Managing Mailboxes: The Essentials

You often need to manage user mailboxes the way you do user accounts. Some of the management tasks are intuitive and others aren't. If you have questions, be sure to read the sections that follow.

Whether you are working with on-premises Exchange or Exchange Online, you can use bulk editing techniques to work with multiple user mailboxes at the same time. To select multiple user mailboxes not in sequence, hold down the Ctrl key and then click the left mouse button on each user mailbox you want to select. To select a series of user mailboxes, select the first mailbox, hold down the Shift key, and then click the last mailbox.

The actions you can perform on multiple resources depend on the types of recipients you've selected. The actions you can perform on multiple user mailboxes include:

- [Updating contact information, organization information, or custom attributes](#)
- [Changing mailbox quotas or deleted item retention settings](#)
- [Enabling or disabling Outlook Web App, POP3, IMAP, MAPI, or ActiveSync](#)
- [Managing policy for Outlook Web App, ActiveSync, Address Books, Retention, Role Assignment, or Sharing](#)
- [Enabling or disabling mailbox archives](#)
- [Moving mailboxes to another database](#)

Although you cannot bulk edit room or equipment mailboxes, you can perform these actions on shared mailboxes.

Viewing Current Mailbox Size, Message Count, and Last Logon

You can use Exchange Admin Center to view the last logon date and time, the mailbox size, and how much of the total mailbox quota has been used by completing these steps:

1. Select **Recipients** in the Navigation menu and then select **Mailboxes**.
2. Double-click the mailbox with which you want to work.
3. On the Mailbox Usage page, review the Last Logon text box to see the last logon date and time (see Figure 7-1). If a user hasn't logged on to her mailbox, you can't get mailbox statistics and will get an error when you view this page.
4. Under the last logon time, notice the mailbox usage statistics, depicted in a bar graph and numerically as a percentage of the total mailbox quota that has been used.

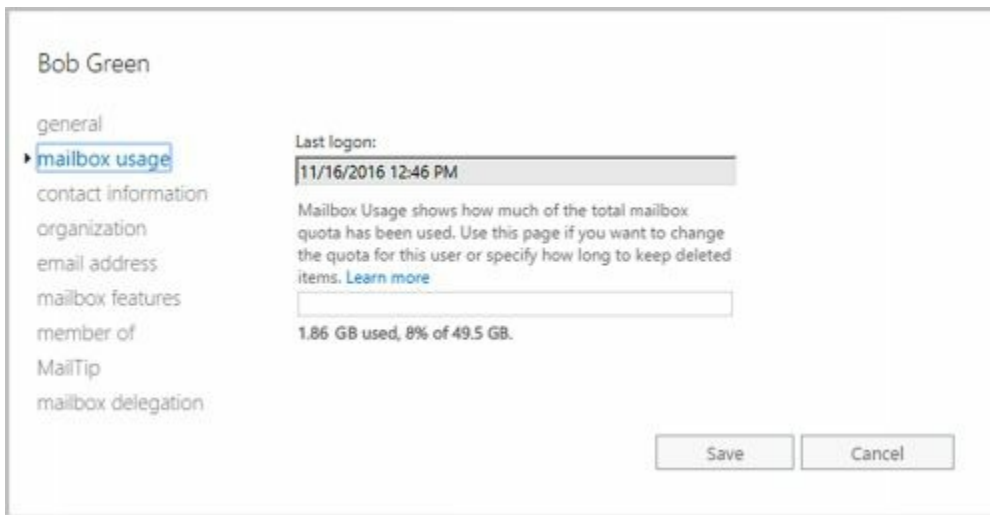


FIGURE 7-1 View mailbox statistics.

If you want to view similar information for all mailboxes on a server, the easiest way is to use the `Get-MailboxStatistics` cmdlet with the `-Server` or `-Database` parameter.

Sample 7-1 shows examples using `Get-MailboxStatistics`. Use the `-Archive` parameter to return mailbox statistics for the archive mailbox associated with a specified mailbox.

SAMPLE 7-1 Getting statistics for multiple mailboxes

```
Get-MailboxStatistics -Identity ' Identity ' [-Archive <$true|$false>]
[-DomainController DomainController ] [-IncludeMoveHistory <$true|$false>]
[-IncludeMoveReport <$true|$false>]
```

```
Get-MailboxStatistics -Server 'Server' | -Database 'Database'
[-DomainController DomainController ]
```

Usage

```
Get-MailboxStatistics -Server ' corpssvr127 '
```

```
Get-MailboxStatistics -Database 'Engineering Primary'
```

```
Get-MailboxStatistics -Identity 'imaginedlands\williams'
```

When you are working with Exchange Management Shell, the standard output won't necessarily provide all the information you are looking for. Often, you need to format the output as a list or table using `Format-List` or `Format-Table`, respectively, to get the additional information you are looking for. `Format-List` is useful when you are working with a small set of resources or want to view all the properties that are available. Once you know what properties are available for a particular resource, you can format the output as a table to view specific properties. For example, if you format the output of `Get-MailboxStatistics` as a list, you see all the properties that are available for mailboxes, as shown in this example and sample output:

```
get-mailboxstatistics -identity "imaginedlands\denises" | format-list
AssociatedItemCount      : 21622
DeletedItemCount        : 1211
```

DisconnectDate :
 DisplayName : Denise Strong
 ItemCount : 20051
 LastLoggedOnUserAccount : NT AUTHORITY\SYSTEM
 LastLogoffTime : 12/21/2016 10:12:33 PM
 LastLogonTime : 12/20/2016 07:36:25 AM
 LegacyDN : /O=FIRST ORGANIZATION/OU=EXCHANGE ADMINISTRATIVE
 GROUP/CN=RECIPIENTS/CN=ERIK ANDERSEN
 MailboxGuid : b7fb0ca8-936b-410f-a2a1-59825eebbdfe
 MailboxType : Private
 ObjectClass : Mailbox
 StorageLimitStatus :
 TotalDeletedItemSize : 1927 KB (1927,535 bytes)
 TotalItemSize : 191121.2 KB (191,121,225 bytes)
 Database : Customer Service Primary
 ServerName : MAILSERVER92
 DatabaseName : Customer Service Primary
 IsQuarantined : False
 IsArchiveMailbox : False
 IsMoveDestination : False
 DatabaseIssueWarningQuota : 1.899 GB (2,039,480,320 bytes)
 DatabaseProhibitSendQuota : 2 GB (2,147,483,648 bytes)
 DatabaseProhibitSendReceiveQuota : 2.3 GB (2,469,396,480 bytes)
 Identity : b7fb0ca8-936b-410f-a2a1-59825eebbdfe
 MapIDentity : b7fb0ca8-936b-410f-a2a1-59825eebbdfe
 OriginatingServer : mailserver92.imaginedlands.com
 IsValid : True
 ObjectState : Unchanged

Once you know the available properties, you can format the output as a table to get exactly the information you want to see. The following example gets information about all the mailboxes in the Engineering Primary database and formats the output as a table:

```
Get-MailboxStatistics -Database 'Engineering Primary' | format-table
DisplayName, TotalItemSize, TotalDeletedItemSize, Database, ServerName
```

Configuring Apps for Mailboxes

With both on-premises Exchange and Exchange Online, you can add apps to the Outlook Web App interface to add functionality. Several apps are installed and made available to users by default, including the following apps created by Microsoft:

- **Action Items** Makes action item suggestions based on message content
- **Bing Maps** Allows users to map addresses found in their messages
- **My Templates** Allows users to save text and images to insert into messages.
- **Suggested Meetings** Shows meeting suggestions found in messages and allows users to add the meetings to their calendars.
- **Unsubscribe** Allows users to block or unsubscribe from email subscription feeds.

Other apps can be added from the Office Store, from a URL, or from a file. All of these apps have various levels of read, read/write, or other permissions on user mailboxes. Because apps also may send data to a third-party service, you may want to consider carefully whether apps should be enabled in your organization. Where strict, high security is a requirement, my recommendation is to disable all apps.

In Exchange Admin Center, you manage apps as part of the organization configuration. Select Organization in the Navigation menu and then select Apps. As shown in Figure 7-2, you'll then see the installed apps and their status. To work with Apps for Outlook, you must have View-Only Organization Management, Help Desk or Organization Management permissions.

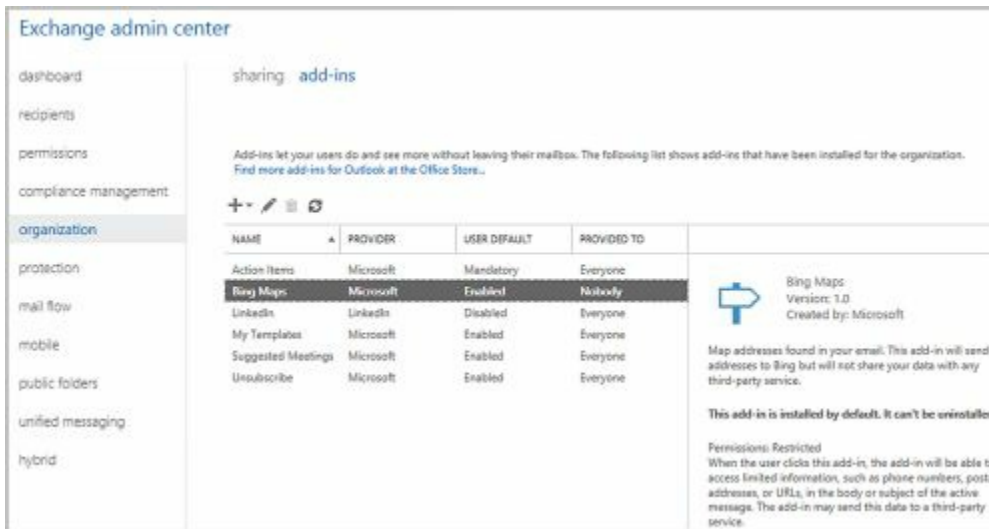



FIGURE 7-2 View the available apps and their status.

To add an app, do one of the following:

- To add an app from the Office store, click New (), select **Add From The Office Store** to open a new browser window to the Office store, and then select an app to add. When you select the app's Add option, review the app details and then click Add. When prompted to confirm, select Yes.
- If you know the URL of the manifest file for the app you want to add, click New and then select Add From URL. In the Add From URL dialog box, enter the URL and then click Install. Be sure to use the full path.
- If you've copied the manifest file to a local server, click New and then select Add From File. In the Add From File dialog box, select Browse. In the Choose File To Upload dialog box, locate and select the manifest file and then select Open. Manifest files end with the .xml extension.

All apps have two status values:

- **User Default** Reflects whether the app is disabled by default, enabled by default, or enabled and mandatory.
- **Provided To** Reflects whether the app is available to all users in the organization (everyone) or to no users in the organization (nobody).

The default apps are made available to all users and enabled by default. This is reflected in the status of Enabled for User Default and Everyone for Provided To by default.

When you install a new app, the app is made available to all users but disabled by default. This is reflected in the status of Disabled for User Default and Everyone for Provided To.

If you have appropriate permissions, you can manage app status by clicking the app and then clicking Edit. In the Action Items dialog box, shown in Figure 7-3, do one of the following:

- If you don't want the app to be available to users, clear the **Make This App Available** checkbox and then click **Save**.
- If you want the app to be available to users, select the **Make This App Available** checkbox and then specify the app status as optional and enabled by default, optional and disabled by default, or mandatory and always enabled. Finally, click **Save**.

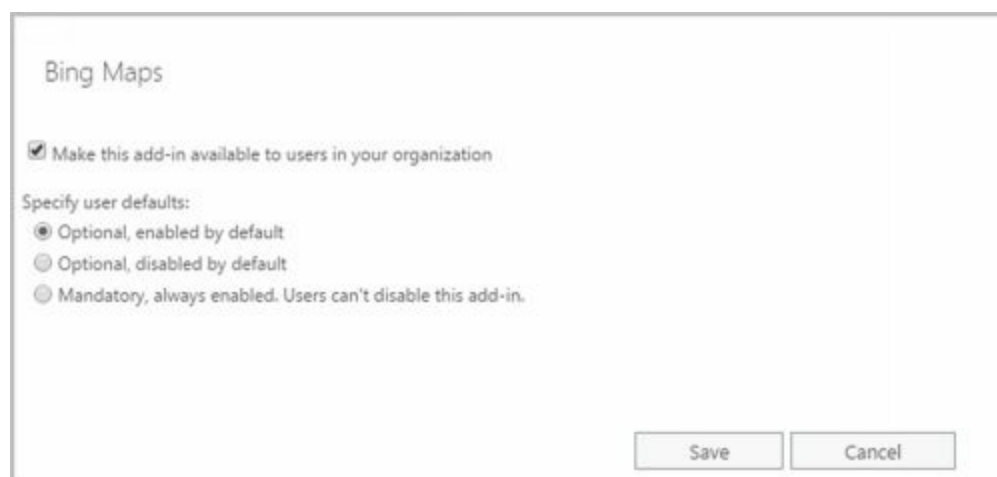


FIGURE 7-3 Manage the app status and availability.

Any app you install can be removed by selecting it and then selecting the Delete option. Although you can't uninstall the defaults apps, you can make any or all of the default apps unavailable to users.

Hiding Mailboxes from Address Lists

Occasionally, you might want to hide a mailbox so that it doesn't appear in the global address list or other address lists. One reason for doing this is if you have administrative mailboxes that you use only for special purposes. To hide a mailbox from the address lists, follow these steps:

1. Open the Properties dialog box for the mailbox-enabled user account by double-clicking the user name in Exchange Admin Center.
2. On the General page, select **Hide From Address Lists**.
3. Click **Save**.

Z William Stanek

general

mailbox usage

contact information

organization

email address

mailbox features

member of

MailTip

mailbox delegation

First name:
Z William

Initials:

Last name:
Stanek

*Name:
William Stanek

*Display name:
Z William Stanek

*Alias:
williams

*User logon name:
williams @ imaginedlands.local

Require password change on next logon

Hide from address lists

[More options...](#)

Save Cancel

Defining Custom Mailbox Attributes for Address Lists

Address lists, such as the global address list, make it easier for users and administrators to find available Exchange resources, including users, contacts, distribution groups, and public folders. The fields available for Exchange resources are based on the type of resource. If you want to add more values that should be displayed or searchable in address lists, such as an employee identification number, you can assign these values as custom attributes.

Exchange provides 15 custom attributes—labeled Customer Attribute 1, Custom Attribute 2, and so on through Custom Attribute 15. You can assign a value to a custom attribute by completing the following steps:

4. Open the Properties dialog box for the mailbox-enabled user account by double-clicking the user name in Exchange Admin Center.

Z William Stanek

general

mailbox usage

contact information

organization

email address

mailbox features

member of

MailTip

mailbox delegation

*User logon name: williams @ imaginedlands.local

Require password change on next logon

Hide from address lists


Organizational unit: imaginedlands.local/Users

Mailbox database: Mailbox Database

Custom attributes:

NUMBER	VALUE

Save Cancel

- On the General page, click **More Options**. Under the Custom Attributes heading, you'll see any currently defined custom attributes. Click **Edit** () to display the Custom Attributes dialog box.
- Enter attribute values in the text boxes provided. Click **OK** and then click **Save**.

Restoring On-Premises Users and Mailboxes

When you disable or delete a mailbox, on-premises Exchange retains the deleted mailbox in the mailbox database and puts the mailbox in a disabled state. There is, however, an important distinction between disabling and deleting a mailbox, and this difference affects recovery. When you disable a mailbox, the Exchange attributes are removed from the user account and the mailbox is marked for removal, but the user account is retained. When you delete a mailbox, the Exchange attributes are removed from the user account, the mailbox is marked for removal, and the user account itself is either marked for deletion or deleted entirely. Additionally, with either, if the mailbox has an in-place archive, the in-place archive will also be marked for removal. However, if the mailbox has a remote archive, the remote archive is removed permanently.

Disabled and deleted mailboxes are referred to as disconnected mailboxes. Disconnected mailboxes are retained in a mailbox database until the deleted mailbox retention period expires, which is 30 days by default. Deleted users may be retained as well.

In Exchange Admin Center, you can find disconnected mailboxes and reconnect them by

completing these steps:

1. Select **Recipients** in the Navigation menu and then select **Mailboxes**.
2. Click the More button (**⋮**) 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 7-4, use the selection list provided to select the server where you want to look for disconnected mailboxes.

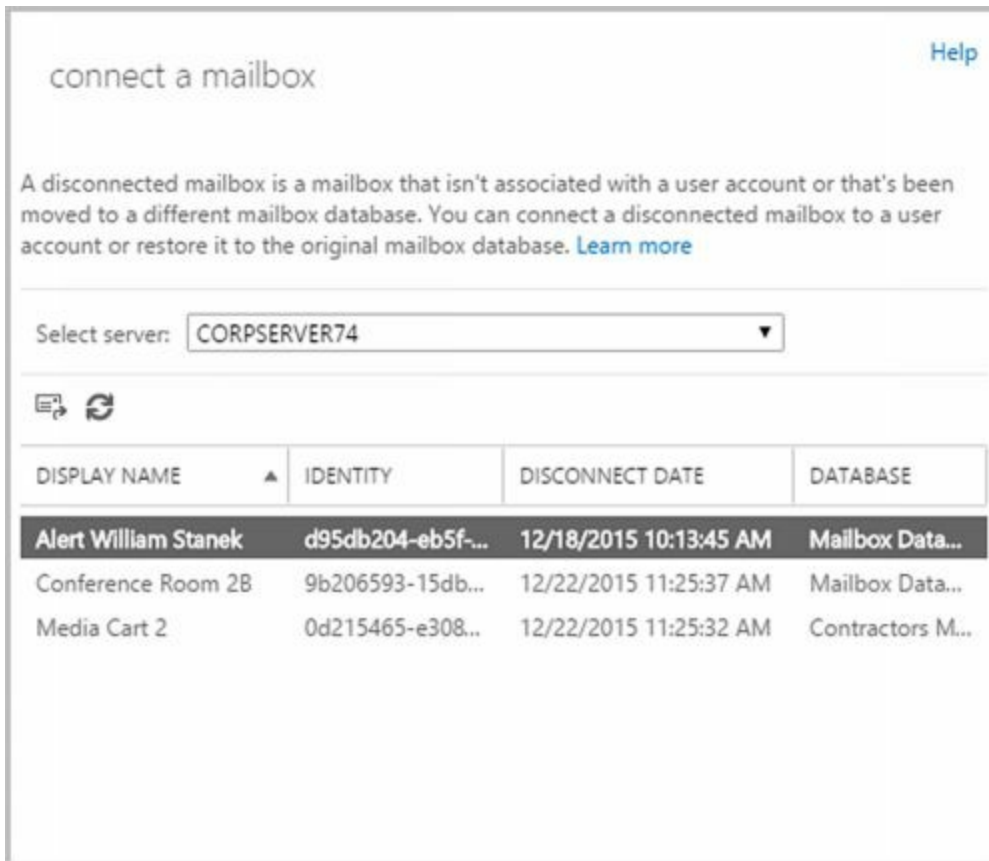



FIGURE 7-4 View disconnected mailboxes.

4. Click the mailbox to restore it and then click Connect ().
5. 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.

IMPORTANT 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 "SoftDeleted" to find soft-deleted mailboxes.

You can find all disabled mailboxes in an on-premises Exchange organization by

entering the following command:

```
Get-MailboxDatabase | Get-MailboxStatistics | Where { $_.DisconnectReason  
-eq "Disabled" } | ft DisplayName,Database,DisconnectDate,DisconnectReason
```

Or you can find disabled mailboxes in a particular database using the following command:

```
Get-MailboxStatistics-Database DatabaseName | Where { $_.DisconnectReason  
-eq "Disabled" } | ft DisplayName,Database,DisconnectDate,DisconnectReason
```

NOTE You can't use this technique with Exchange Online. See "Restoring online users and mailboxes" later in this chapter.

If you find that you need a mail-enabled or mailbox user account that was deleted, you may be able to restore the deleted account. For on-premises Exchange, you can restore user accounts from Active Directory. When Active Directory Recycle Bin is enabled, you can recover deleted objects using Active Directory Administrative Center (as long as the deleted object and recycled object lifetimes have not expired).

In Active Directory Administrative Center, select the Deleted Object container to see the available deleted objects. When you select a deleted user by clicking it, you can use the Restore option to restore the user to its original container. For example, if the user account was deleted from the Users container, the user account is restored to this container. Once the user account is restored, you can restore the Exchange settings and data. You can use Connect-Mailbox to connect the user account to its disconnected mailbox.

When you connect a disconnected mailbox using Connect-Mailbox, you associate the mailbox with a user account that isn't mail-enabled, which means the user account cannot have an existing mailbox associated with it. Connect-Mailbox has a slightly different syntax for standard mailboxes, shared mailboxes, and linked mailboxes. For standard mailbox users, the basic syntax for Connect-Mailbox is:

```
Connect-Mailbox -Identity ExchangeId -Database DatabaseName -User ADUserId  
-Alias ExchangeAlias
```

where ExchangeID identifies the disconnected mailbox in the Exchange organization, DatabaseName is the name of the database where the disconnected mailbox resides, ADUserID identifies the Active Directory user account to reconnect the mailbox to, and ExchangeAlias sets the desired Exchange Alias. Consider the following example:

```
Connect-Mailbox -Identity "Thomas Axen" -Database "Sales Database"  
-User "Thomas Axen" -Alias ThomasA
```

This example reconnects the Exchange mailbox for Thomas Axen with the related user account in Active Directory and sets the Exchange alias as ThomasA. The alias is combined with the user logon domain to set the User Principal Name (referred to in the UI as the User Logon Name). The User Principal Name must be unique within the organization. If another user account has the same User Principal Name, you'll see a

warning about a user name conflict. You will need to resolve this conflict before you can connect the mailbox.

When you disable or remove an archive mailbox from a mailbox, the archive mailbox is disconnected from the source mailbox, marked for deletion, and retained according to the retention settings. To connect a disabled archive mailbox to the original source mailbox, you use the Connect-Mailbox cmdlet with the -Archive parameter.

Although Connect-Mailbox has restrictions, you can connect a disconnected mailbox to a user account that already has a mailbox. When you restore the mailbox, its contents are copied into the target user's existing mailbox while the deleted mailbox itself is retained in the mailbox database until the retention period expires (or it is purged by an administrator).

You use New-MailboxRestoreRequest to restore mailboxes to accounts with existing mailboxes. The basic syntax is:

```
New-MailboxRestoreRequest-SourceMailbox MailboxID -SourceDatabase  
DatabaseName -TargetMailbox ExchangeID
```

where MailboxID is the display name or GUID of the disconnected mailbox to restore, DatabaseName is the name of the database where the disconnected mailbox resides, and ExchangeID is an Exchange alias or name for the account where the mailbox should be added. Consider the following example:

```
New-MailboxRestoreRequest -SourceMailbox "Karen Berg" -SourceDatabase  
"Marketing Database" -TargetMailbox "Dag Rovik"
```

You can restore archive mailboxes to users with existing accounts as well. Use the -TargetIsArchive parameter as shown in this example:

```
New-MailboxRestoreRequest -SourceMailbox "In-Place Archive - Karen Berg"  
-SourceDatabase "Marketing Database" -TargetMailbox "Dag Rovik"  
-TargetIsArchive
```

Restoring Online Users and Mailboxes

If you remove the Exchange Online license for an online user account, the user's account is marked as an unlicensed account. Exchange Online deletes mailboxes from unlicensed accounts automatically after the grace period expires. By default, this grace period is 30 days. If you delete a user account in the online organization, the user account is marked as deleted but retained until the retention period expires, which is 30 days by default.

In Office 365 Admin Center, you can find deleted users and restore them by completing these steps:

1. Select Deleted Users under the Users heading in the Navigation menu to view deleted users, as shown in Figure 7-5. If the online organization has available licenses, you can restore the deleted users.



FIGURE 7-5 View deleted but retained users in Office 365 Admin Center.

2. Select the account to restore and then click **Restore**. As shown in Figure 7-6, you'll then be prompted to confirm the action by clicking Restore again.

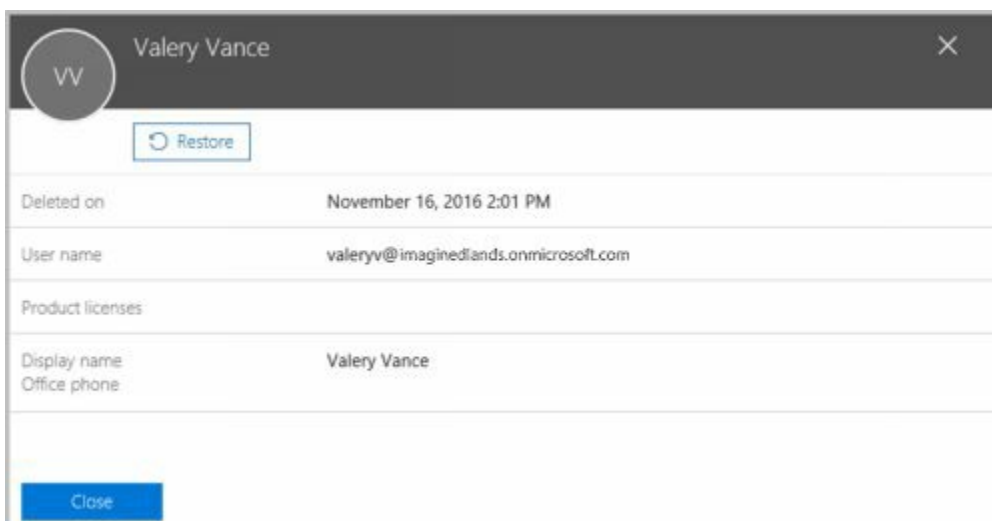
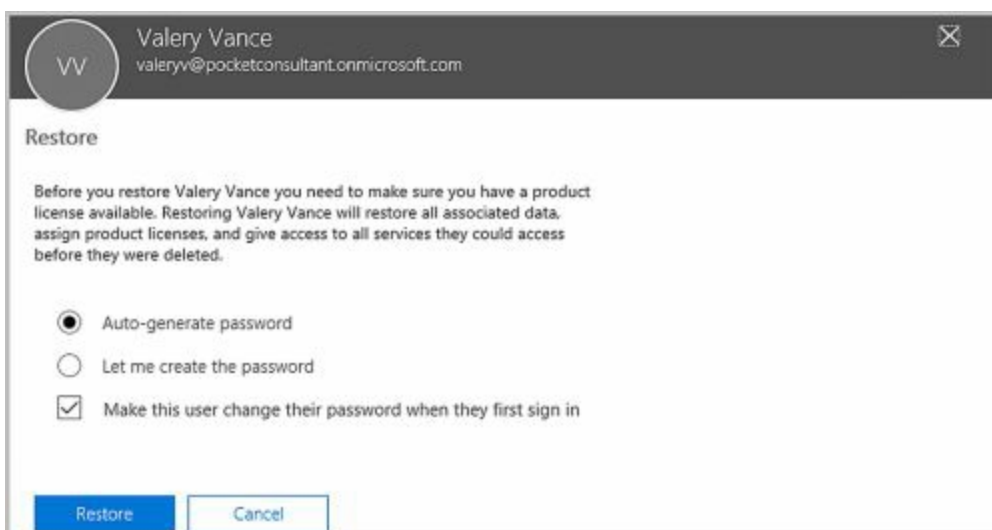


FIGURE 7-6 Restore online users in Office 365 Admin Center.

3. Before you confirm the action by clicking Restore again when prompted, consider whether you want to assign the user an automatically generated password or one you designate. You also can specify whether the user must change their password when they first sign on.



4. Whether you assign a password or generate one, the new password is sent in email to a designated administrator. You can change the default as necessary or add recipients. Be sure to separate each email address with a semicolon. When you are ready to continue, click **Send Email And Close**.
5. Repeat Steps 2 – 4 to restore other deleted accounts.

NOTE Keep in mind that account restoration will fail if there are any naming or other conflicts. The User Principal Name must be unique within the organization. If another user account has the same the User Principal Name, you'll see a warning about a user name conflict. You'll then be able to edit the user name or replace the active user with the deleted user.

When you connect to Microsoft Online Services as discussed in Chapter 2 "Working with Exchange Online," you can get information about accounts in Windows PowerShell. Enter **Get-MsolUser** to get a list of active user accounts. As shown in the following example the default output shows the User Principal Name, display name, and licensing status of user accounts:

UserPrincipalName	DisplayName	isLicensed
cart3@imaginedlands.onmicrosoft.com	Media Cart 3	False
wrstanek@imaginedlands.onmicrosoft.com	William Stanek	True
room3@imaginedlands.onmicrosoft.com	Conference Room 3	False
georges@imaginedlands.onmicrosoft.com	George Schaller	False
room42@imaginedlands.onmicrosoft.com	Conference Room 42	False

The output shows the user accounts associated with all types of users, including the user accounts associated with room and equipment mailboxes. Although room and equipment mailboxes don't need to be licensed, standard user accounts require licenses.

You can get a list of users whose accounts have been marked for deletion by entering **Get-MsolUser -ReturnDeletedUsers**. Accounts marked for deletion are listed by User Principal Name, display name, and licensing status. To restore a deleted account, use **Restore-MsolUser**. The basic syntax for this command is:

```
Restore-MsolUser -UserPrincipalName OnlineId
```

where **OnlineId** is the User Principal Name of the account to restore:

```
Restore-MsolUser -UserPrincipalName valv@imaginedlands.onmicrosoft.com
```

The account restore will fail if there are any naming or other conflicts. To resolve a name conflict, use the **-NewUserPrincipalName** parameter to set a new User Principal Name for the user.

Repairing Mailboxes

You can use **New-MailboxRepairRequest** to detect and repair mailbox corruption. By default, the command attempts to repair all types of mailbox corruption issues, including issues associated with search folders, aggregate counts, provisioned folders, and folder views.

The basic syntax for **New-MailboxRepairRequest** is:

```
New-MailboxRepairRequest -Mailbox ExchangeID
```

where **ExchangeID** identifies the mailbox to repair, such as:

New-MailboxRepairRequest -Mailbox TonyS

New-MailboxRepairRequest -Mailbox tonys@imaginedlands.com

New-MailboxRepairRequest -Mailbox "Tony Smith"

During the repair process, the mailbox cannot be accessed. Once started, the detect and repair process cannot be stopped, unless you dismount the associated database. Add the -Archive parameter to repair the archive mailbox associated with an Exchange identifier rather than the primary mailbox.

You also can use New-MailboxRepairRequest to examine and repair all mailboxes in a database. As the repair process works its way through all the mailboxes in the database, only the mailbox being repaired is locked and inaccessible. All other mailboxes in the database remain accessible to users.

Moving Mailboxes

Exchange Server 2016 supports online mailbox moves. To complete an upgrade, balance the server load, manage drive space, or relocate mailboxes, you can move mailboxes from one server or database to another server or database. The process you use to move mailboxes depends on where the mailbox or mail data is stored:

- When you want to work with mail data stored on a user's computer, you can use the import or export process to move mail data.
- When a user's mailbox is stored on an on-premises Exchange server and you want to move the mailbox to a database on the same server or another server in the same forest, you can use an online mailbox move or batch migration to move the mailbox.
- When a user's mailbox is stored on an on-premises Exchange server in one Active Directory forest and you want to move the mailbox to an on-premises Exchange server in another forest, you can use a cross-forest move to move the mailbox.
- When a user's mailbox is stored on-premises and you want to move the mailbox to Exchange Online or vice versa, you can use a remote move to move the mailbox.

Importing and Exporting Mail Data

When Microsoft Outlook uses Exchange Server, a user's mail data can be delivered in one of two ways:

- Server mailbox with local copies
- Personal folders

With server mailboxes, messages are delivered to mailboxes on the Exchange server and users can view or receive new mail only when they are connected to Exchange. A local copy of the user's mail data is stored in an .ost file on her computer.

Personal folders are alternatives to server mailboxes. Personal folders are stored in a .pst file on the user's computer. With personal folders, you can specify that mail should be delivered to the user's inbox and stored on the server or that mail should be delivered only to the user's inbox. Users have personal folders when Outlook is configured to use Internet email or other email servers. Users might also have personal folders if the auto-archive feature is used to archive messages.

When you are working with on-premises Exchange, you can:

- Import mail data from .pst files using mailbox import request cmdlets
- Export mail data to .pst files using mailbox export request cmdlets

IMPORTANT You must have the Mailbox Import Export role to be able to import or export mailbox data. As this role isn't assigned to any role group, you must be explicitly assigned this role.

The import and export processes are asynchronous. They are queued and processed independently of Exchange Management Shell. The related commands are shown in the

following list:

IMPORT MAILBOX DATA	EXPORT MAILBOX DATA
Get-MailboxImportRequest	Get-MailboxExportRequest
New-MailboxImportRequest	New-MailboxExportRequest
Set-MailboxImportRequest	Set-MailboxExportRequest
Suspend-MailboxImportRequest	Suspend-MailboxExportRequest
Resume-MailboxImportRequest	Resume-MailboxExportRequest
Remove-MailboxImportRequest	Remove-MailboxExportRequest
Get-MailboxImportRequestStatistics	Get-MailboxExportRequestStatistics

Mailbox imports and exports are initiated with Mailbox Import and Mailbox Export requests respectively. These requests are sent to the Microsoft Exchange Mailbox Replication Service (MRS) running on a Mailbox server in the source forest. The MRS queues the request for processing, handling all requests on a first-in, first-out basis. When a request is at the top of the queue, the replication service begins importing or exporting mail data.

Before you can import or export data, you need to create a shared network folder that is accessible to your Exchange servers, and the Exchange Trusted Subsystem group must have read/write access to this share.

You use New-MailboxImportRequest to import data from a .pst file to a mailbox or personal archive. Keep in mind you can't import data to a user account that doesn't have a mailbox and that the destination mailbox must be already available. The import process will not create a mailbox. By default, all mail folders are imported. However, you can specifically include or exclude folders. You also can import mail data to only the user's personal archive.

You use New-MailboxExportRequest to export mailbox data to a .pst file. The command allows you to export one or more mailboxes, with each mailbox export handling a separate request. When exporting mail data, you can specify folders to include or exclude and export mail data from the user's archive. You also can filter the messages so only messages that match your content filter are exported.

Performing On-Premises Mailboxes Moves and Migrations

The destination database for an on-premises mailbox move can be on the same server, on a different server, in a different domain, or in a different Active Directory site.

Exchange Server 2016 performs move operations as a series of steps that allows a mailbox to remain available to a user while the move operation is being completed. When the move is completed, the user begins accessing the mailbox in the new location. Because users can continue to access their email account during the move, you can perform online moves at any time.

The online move process hasn't changed substantially since it was introduced with Exchange Server 2010:

- On-premises mailbox moves are initiated with a Move Mailbox request that is sent to the Microsoft Exchange MRS running on a Mailbox server in the source forest. The MRS queues the request for processing, handling all requests on a first-in, first-out basis. When a request is at the top of the queue, the replication service begins replicating mailbox data to the destination database.
- When the replication service finishes its initial replication of a mailbox, it marks the mailbox as Ready To Complete and periodically performs data synchronization between the source and destination database to ensure that the contents of a mailbox are up to date. After a mailbox has been moved, you can complete the move request and finalize the move.

When you are working with PowerShell, you initiate a move using `New-MoveRequest` and then start the actual move using `Start-MoveRequest`. Although the online move process allows you to move multiple mailboxes, with each move handled as a separate request, the process isn't ideal for batch moves of multiple mailboxes, and this is where mailbox migrations come in. With mailbox migration, you can move multiple mailboxes in an Exchange on-premises organization, migrate on-premises mailboxes to Exchange Online, or migrate Exchange Online mailboxes back to an on-premises Exchange organization.

NOTE You can use the batch migration process to move a single or multiple mailboxes within on-premises Exchange. With a single mailbox, the batch migration is handled as a local move.

From a high level, the standard batch migration process is similar to a mailbox move:

- Batch mailbox migration is initiated with a Migration Batch request that is sent to the Microsoft Exchange MRS running on a Mailbox server in the source forest. The MRS queues the request for processing, handling all requests on a first-in, first-out basis. When a request is at the top of the queue, the replication service begins replicating mailbox data to the destination database.
- When the replication service finishes its initial replication of a mailbox, it marks the mailbox as Ready To Complete and periodically performs data synchronization between the source and destination database to ensure that the contents of a mailbox are up to date. After a mailbox has been migrated, you can complete the migration request and finalize the migration.

Where things get complicated are on cross-forest batch migrations and remote migrations. With a cross-forest migration, you perform a batch mailbox migration from

an Exchange server in one Active Directory forest to an Exchange server in another Active Directory forest. With a remote migration, you perform a batch mailbox migration from on-premises Exchange to Exchange Online or vice versa.

Cross-forest and remote migrations use migration endpoints. You create a migration endpoint in the target environment. The endpoint identifies the source environment where the mailboxes are currently located. You then initiate the migration in the target environment. With a cross-forest migration, this means you:

1. Create a migration endpoint in the target domain.
2. Initiate the migration in the target domain.

With a migration from on-premises Exchange to Exchange Online, this means you:



1. Create a migration endpoint in Exchange Online.
2. Initiate the migration from Exchange Online.

With a migration from Exchange Online to on-premises Exchange, this means you:

1. Create a migration endpoint in on-premises Exchange.
2. Initiate the migration from on-premises Exchange.

A complete cross-forest or remote migration has four parts. You create a migration endpoint using `New-MigrationEndpoint` and then create the migration batch using `New-MigrationBatch`. You start the migration using `Start-MigrationBatch`. When the migration has finished initial synchronization, you can finalize the migration using `Complete-MigrationBatch`.

In Exchange Admin Center, you can initiate move and migration requests using the options on the Migration page. To access this page, select **Recipients** in the Navigation menu and then select **Migration** (see Figure 7-7). Although the PowerShell commands for moves and migrations give you complete control over the process, you'll find that Exchange Admin Center greatly simplifies the process:

- For local moves, you log on to a Mailbox server in the Active Directory forest where the source mailboxes are located. On the Migration page, select **New** () and then select **Move To A Different Database**. Follow the prompts in the New Local Mailbox Move dialog box to perform the move.
- For remote migrations, you can use the options in Exchange Admin Center for Exchange Online to initiate the process, whether migrating from or to Exchange Online. On the Migration page, select **More** (), select **Migration Endpoints**, and then follow the prompts to create the required migration endpoint. Next, select **New** and then select either **Migrate To Exchange Online** or **Migrate From Exchange Online** as appropriate. Follow the prompts in the New Migration Batch dialog box to perform the migration.
- For cross-forest moves, you log on to a Mailbox server in the target Active Directory

forest. On the Migration page, select **More** (**⋮**), select **Migration Endpoints**, and then follow the prompts to create the required migration endpoint. Next, select **New** (**+**) and then select **Move To This Forest**. Follow the prompts in the New Cross-Forest Mailbox Move dialog box to perform the move.

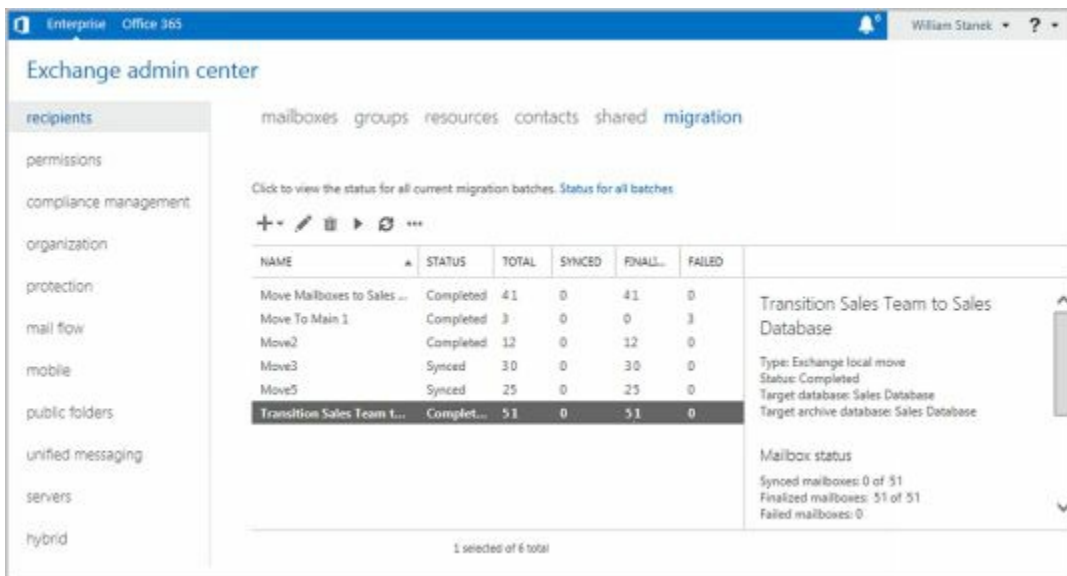


FIGURE 7-7 Check the status of move and migration requests.

On the Migration page, you also can track the status of move and migration requests. If a move or migration request fails, you can get more information about the failure by double-clicking the request and then clicking View to the right of the Failed Message entry.

When you move mailboxes from one server to another, to a different organization, or even to a different database on the same server, keep in mind that the Exchange policies of the new mailbox database might be different from the old one. Because of this, consider the following issues before you move mailboxes to a new server or database:

- **General policy** Changes to watch out for include the storage limits, the deleted item retention, and the default offline address book settings. The risk is that the users whose mailboxes you move could lose or gain access to public folders. They might have a different offline address book, which might have different entries. This address book will also have to be downloaded in its entirety the first time the user's mail client connects to Exchange after the move.
- **Database policy** Changes to watch out for pertain to the maintenance interval and automatic mounting. If Exchange performs maintenance when these users are accessing their mail, they might have slower response times. If the mailbox database is configured so that it isn't mounted at startup, restarting the Exchange services could result in the users not being able to access their mailboxes.
- **Limits** Changes to watch out for pertain to storage limits and deletion settings. Users might be prohibited from sending and receiving mail if their mailbox exceeds the storage limits of the new mailbox database. Users might notice that deleted items stay in their Deleted Items folder longer or are deleted sooner than expected if the

Keep Deleted Items setting is different.

Performing On-Premises Mailbox Moves

With online moves and batch migrations, you can move mailboxes between databases on the same server. You also can move mailboxes from a database on one server to a database on another server regardless of whether the servers are in a different Active Directory site or in another Active Directory forest.

Normally, when you perform online migrations, the move process looks like this:

1. You create a batch migration request for the mailboxes that you want to move using either Exchange Admin Center or Exchange Management Shell.
2. The request is sent to the Mailbox Replication Service running on a Mailbox server in the current Active Directory site. This server acts as the Mailbox Replication Service proxy.
3. MRS adds the mailboxes to the Request queue and assigns the status Created to the request. This indicates the move has been requested but not started.
4. When a request is at the top of the queue, MRS begins replicating the related mailboxes to the destination database and assigns the Syncing status to the request.
5. When MRS finishes its initial replication of the mailboxes, the service assigns the Synced status to the request.
6. The request remains in the Synced state until you or another administrator specifies that you want to complete the request. MRS performs a final data synchronization and then marks the request as Completed.
7. When the request is completed, the mailboxes are available in the new location. Because users can continue to access their email accounts during the move, you can perform online moves and migrations at any time.

One way to perform online mailbox moves and migrations is by using Exchange Management Shell. The commands for performing online mailbox moves include the following:

- **Get-MoveRequest** Displays the detailed status of an ongoing mailbox move that was initiated using the New-MoveRequest cmdlet.
- **New-MoveRequest** Starts a mailbox move. You also can verify readiness to move by using the -WhatIf parameter. Use the -Priority parameter to set the relative priority of the request.
- **Resume-MoveRequest** Resumes a move request that has been suspended or failed.
- **Set-MoveRequest** Changes a move request after it has been started.
- **Suspend-MoveRequest** Suspends a move request that has been started but has not yet been completed.
- **Remove-MoveRequest** Cancels a mailbox move initiated using the New-MoveRequest cmdlet. You can use the Remove-MoveRequest command any time after

initiating the move but only if the move request is not yet complete.

The commands for performing batch mailbox migrations include the following:

- **Get-MigrationBatch** Displays the detailed status of an ongoing mailbox migration that was initiated using the New-MigrationBatch cmdlet.
- **Set-MigrationBatch** Changes a migration request after it has been started.
- **New-MigrationBatch** Submits a new mailbox migration request. You also can verify readiness to migrate by using the -WhatIf parameter. Use the -AutoStart parameter to allow immediate processing of the request. Use the -AutoComplete parameter to automatically finalize the batch when the initial synchronization is complete.
- **Start-MigrationBatch** Submits a migration request for processing; required when the -AutoStart parameter is not used with New-MigrationBatch.
- **Stop-MigrationBatch** Stops a migration request that has been started but has not yet been completed.
- **Complete-MigrationBatch** Finalizes a migration request that has been synchronized; required when the -AutoComplete parameter is not used with New-MigrationBatch.
- **Remove-MigrationBatch** Deletes a mailbox migration request that either isn't running or has been completed. If you created a new request but haven't submitted it, you can use this command to remove the request so that the mailboxes specified in the request aren't migrated. If the request is completed, the mailboxes are already migrated and you can use this command to remove the request from the queue.
- **Get-MigrationUser** Retrieves information about the ongoing migration of a particular mailbox.
- **Remove-MigrationUser** Allows you to remove a mailbox from a migration request.
- **Test-MigrationServerAvailability** Ensures the target server for a cross-premises move is available and verifies the connection settings.

Other batch migration commands include: Get-MigrationStatistics, Get-MigrationUserStatistics, Get-MigrationConfig, Set-MigrationConfig, Get-MigrationEndpoint, Set-MigrationEndpoint, New-MigrationEndpoint, and Remove-MigrationEndpoint.

Moving Mailboxes Within a Single Forest

You perform online mailbox moves within a single forest by using Exchange Management Shell. To verify move readiness, use New-MoveRequest with the -WhatIf parameter for each mailbox you plan to move. The following examples show two different ways you can verify whether Morgan Skinner's mailbox can be moved:

```
New-MoveRequest -Identity 'morgans'  
-TargetDatabase "Engineering Primary" -WhatIf
```

'imaginedlands.com/users/Morgan Skinner' | New-MoveRequest

-TargetDatabase 'Engineering Primary' -WhatIf

To initiate an online move, you use New-MoveRequest for each mailbox you want to move. The following examples show two different ways you can move Morgan Skinner's mailbox:

```
New-MoveRequest -Identity 'morgans' -Remote -RemoteHostName  
'mailserver62.imaginedlands.com' -mrserver  
'mailserver19.imaginedlands.com' -TargetDatabase "Engineering Primary"
```

```
'imaginedlands.com/users/Morgan Skinner' | New-MoveRequest -Remote  
-RemoteHostName 'mailserver62.imaginedlands.com' -mrserver  
'mailserver19.imaginedlands.com' -TargetDatabase 'Engineering Primary'
```

After you initiate a move, you can check the status of the online move using Get-MoveRequest. As shown in the following example, the key parameter to provide is the identity of the mailbox you want to check:



```
Get-MoveRequest -Identity 'morgans'
```

You can use Suspend-MoveRequest to suspend a move request that has not yet completed, and Resume-MoveRequest to resume a suspended move request. Resuming a suspended request allows it to complete.

You can cancel a move at any time prior to running the move request being completed by Exchange. To do this, run Remove-MoveRequest and specify the identity of the mailbox that shouldn't be moved. An example follows:

```
Remove-MoveRequest -Identity 'morgans'
```

When your source and destination Mailbox servers are running Exchange Server 2016 and are in the same forest, you can move mailboxes by completing these steps:

1. Log on to Exchange Admin Center via a Mailbox server in the domain or forest you want to work with. In Exchange Admin Center, select Recipients in the Navigation menu and then select Migration.
2. On the Migration page, you select New () and then select **Move To A Different Database**. This starts the New Local Mailbox Move Wizard.
3. On the Select The Users page, shown in Figure 7-8, you can select the mailboxes to migrate by doing one of the following:
 - Select the mailboxes that you want to migrate using the graphic interface. Click Add (). Use the Select Mailbox dialog box to select the mailboxes to move and then click Add. Next, click OK.

You can select and move multiple mailboxes at the same time. To select multiple mailboxes individually, hold down the Ctrl key, and then click each mailbox that you want to select. To select a sequence of mailboxes, select the first mailbox, hold

down the Shift key, and then click the last user mailbox.

- Select the mailboxes that you want to migrate using a file containing a list of comma-separated Exchange identifiers. Click **Specify The Users With A CSV File** and then click **Choose File**. Use the Open dialog box to select the .csv file and then click **OK**.

The file you use should be named with the .CSV extension. The first line of the file should identify the column of data to import as: `EmailAddress` and each successive line in the file should be the email address of a mailbox to migrate, as shown in this example:

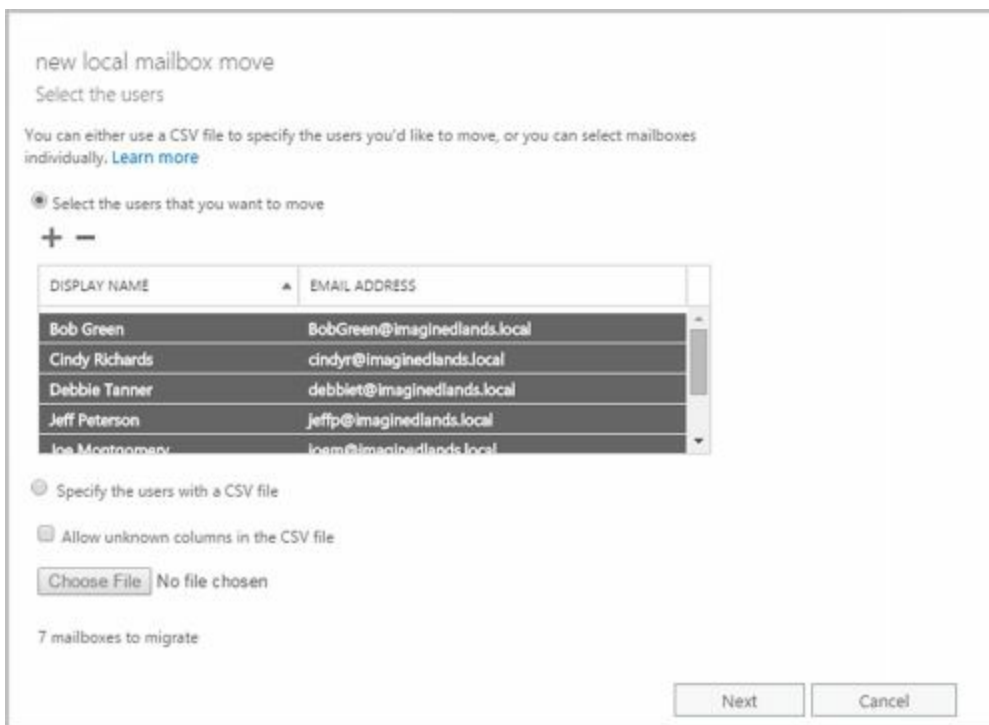


FIGURE 7-8 Select the mailboxes to migrate.

EmailAddress
annal@imaginedlands.com
deanh@imaginedlands.com
indron@imaginedlands.com
paulab@imaginedlands.com
williams@imaginedlands.com

4. Click **Next**. On the Move Configuration page, shown in Figure 7-9, enter a descriptive name for the migration batch.

FIGURE 7-9 Configure the settings for the move request.

5. Use the Archive options to specify whether you want to move only the primary mailbox for the selected recipients, only the archive mailbox for the selected recipients, or both.
6. If you are moving the primary mailboxes for recipients, click **Browse** to the right of the Target Database text box. In the Select Mailbox Database dialog box, choose the mailbox database to which the mailbox should be moved. Mailbox databases are listed by name as well as by associated server and Exchange version.
7. If you are moving the archive mailboxes for recipients, click **Browse** to the right of the Target Archive Database text box. In the Select Mailbox Database dialog box, choose the mailbox database to which the mailbox should be moved. Mailbox databases are listed by name as well as by associated server and Exchange version.
8. If corrupted messages are found in a mailbox that you are migrating, the messages are skipped automatically and not migrated as part of the mailbox. By default, the wizard skips a limited number of bad items in each mailbox and stops the migration if this value is exceeded. To specify the maximum number of bad items that can be skipped in each mailbox, enter a new value in the Bad Item Limit text box or enter 0 to allow an unlimited number of bad items to be skipped.
9. Click **Next**. On the Start The Batch page, your current login is selected as the recipient for the batch report. This report will contain details about errors encountered during the migration. To add or change recipients for this report, click **Browse**. Then use the Select Members dialog box to select the recipients that should receive the report and then click **OK**. You must select at least one recipient.

new local mailbox move
Start the batch

A new migration batch will be created after you click new. [Learn more](#)

*After the batch is complete, a report will be sent to the following recipients. You must select at least one recipient to receive this report.

William Stanek

Please select the preferred option to start the batch:

Manually start the batch later (by selecting it in the migration dashboard and then clicking Start)

Automatically start the batch

Please select the preferred option to complete the batch:

Manual Complete the batch (by clicking the "Complete this migration batch" link on the right pane, after the link becomes active)

Automatically complete the migration batch

10. By default, Exchange Server creates and starts the batch migration request. When the request is completed, Exchange Server will also automatically finalize it. If you want to manually start the batch, select the Manual Start option. If you want to manually finalize the batch, select the Manual Complete option.
11. Click **New**. Migrating mailboxes can take several hours, depending on the size of the mailboxes you are moving. You can check the status of move requests by refreshing the view on the Migration page. While the request is in the Synced state, you can cancel the request by selecting it and then clicking Delete. You cannot cancel a request that has started syncing.

Moving Mailboxes Between Forests

You can perform online mailbox moves between different Exchange forests using Exchange Admin Center or Exchange Management Shell. When you are moving mailboxes between forests, verify that mailboxes are ready to be moved before you submit a move request. To verify readiness, the Microsoft Exchange Mailbox Replication service proxy in the source forest checks the status of each mailbox you are moving and also ensures you have the permissions required to move the mailboxes from the source forest to the target forest. If a user has an archive mailbox or subscriptions, you will likely need to remove the archive mailbox, the subscriptions, or both before you are able to move the mailbox.


You can verify move readiness in Exchange Management Shell by using `New-MoveRequest` with the `-WhatIf` parameter for each mailbox you plan to move. The following examples show two different ways you can verify whether Rob Carson's mailbox can be moved:

```
New-MoveRequest -Identity 'robc' -Remote
-RemoteHost 'mailserver62.imaginedlands.com' -mrserver
'mailserver19.imaginedlands.com'
-TargetDatabase "Engineering Primary" -WhatIf

'imaginedlands.com/users/Rob Carson' | New-MoveRequest -Remote
-RemoteHost 'mailserver62.imaginedlands.com' -mrserver
```

'mailserver19.imaginedlands.com'
-TargetDatabase 'Engineering Primary' -WhatIf

You can perform online mailbox moves between forests by following these steps:

1. Log on to Exchange Admin Center via a Mailbox server in the target forest. In Exchange Admin Center, select **Recipients** in the Navigation menu and then select **Migration**.
2. On the Migration page, select New () and then select **Move To This Forest**. This starts the New Cross-Forest Mailbox Move Wizard.
3. On the Select The Users page, shown in Figure 7-10, you can select the mailboxes to migrate by doing one of the following:
 - Select the mailboxes that you want to migrate using the graphic interface. Click the Add button. Use the Select Mailbox dialog box to select the mailboxes to move and then click Add. Next, click OK.

You can select and move multiple mailboxes at the same time. To select multiple mailboxes individually, hold down the Ctrl key, and then click each mailbox that you want to select. To select a sequence of mailboxes, select the first mailbox, hold down the Shift key, and then click the last user mailbox.

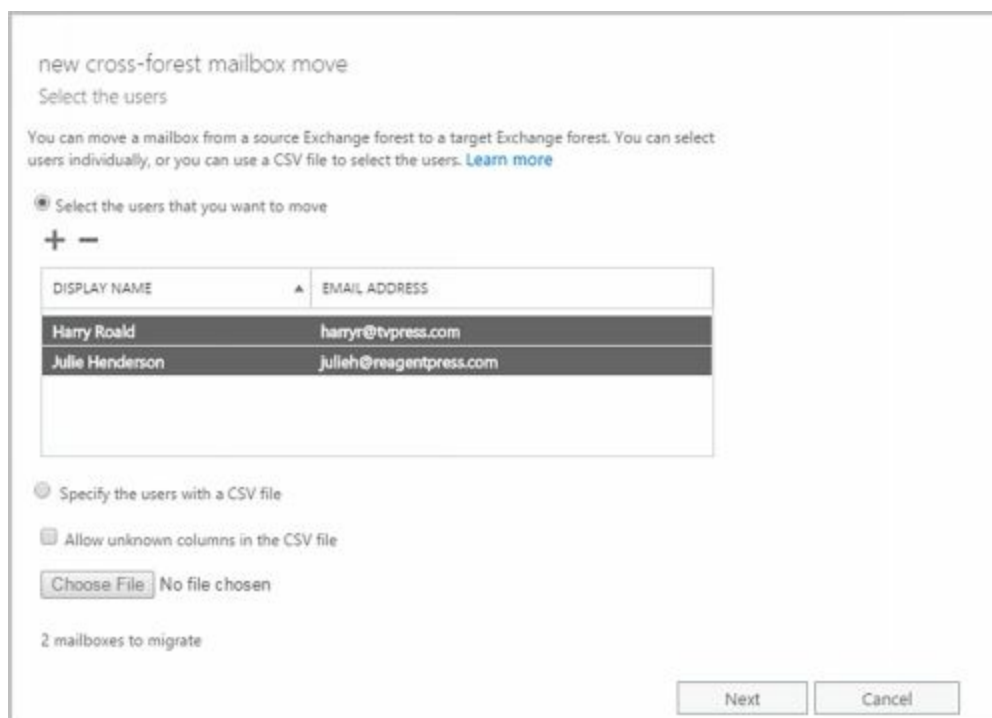


FIGURE 7-10 Configure the settings for the cross-forest move request.

- Select the mailboxes that you want to migrate using a file containing a list of comma-separated Exchange identifiers. Click **Specify The Users With A CSV File** and then click **Choose File**. Use the Open dialog box to select the .csv file and then click **OK**.

The file you use should be named with the .CSV extension. The first line of the file should identify the column of data to import as: EmailAddress and each successive line in the file should be the email address of a mailbox to migrate, as shown in

this example:

EmailAddress

annal@imaginedlands.com

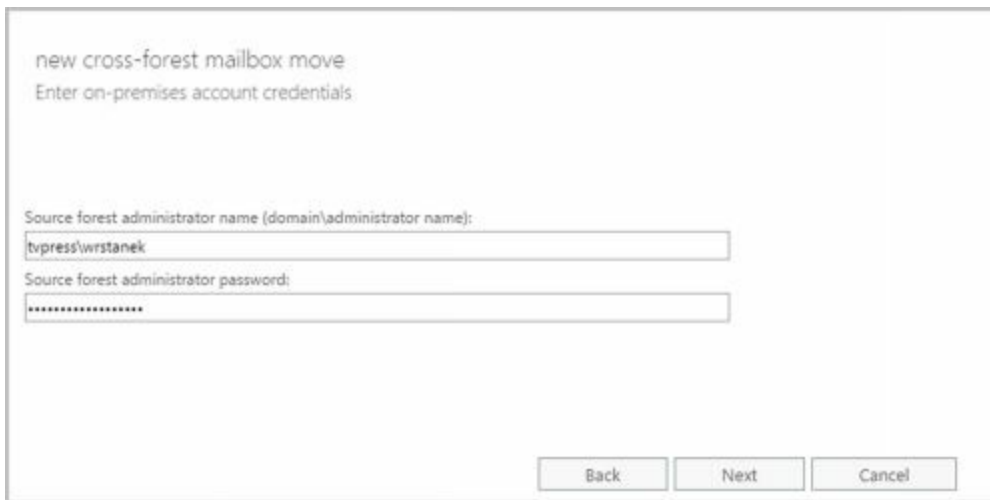
deanh@imaginedlands.com

indron@imaginedlands.com

paulab@imaginedlands.com

williams@imaginedlands.com

4. Click **Next**. The target forest is the forest to which you are connected. The source forest is the forest where the mailboxes are located currently. In the Source Forest Administrator Name text box, enter the name of a user account that has appropriate administrative privileges in the source forest. Enter the name in Domain\UserName format, such as Imaginedlands\Williams.



new cross-forest mailbox move
Enter on-premises account credentials

Source forest administrator name (domain/administrator name):
tvpres\wrstane

Source forest administrator password:

Back Next Cancel

NOTE The administrator must have sufficient permissions to create the required migration endpoint and move accounts. Typically, this means the account must be a member of both the Recipient Management and Server Management groups in the Exchange organization or have Organization Management permissions. However, if you previously migrated accounts between these forests, the migration endpoint created previously may still be available, in which case only Recipient Management permissions are required.

5. In the Source Forest Administrator Password text box, enter the password for the previously specified account.
6. When you click **Next**, Exchange uses the Autodiscover service to try to detect the availability of the migration endpoint as well as to test connectivity. If errors occur, the Confirm The Migration Endpoint page is displayed. At this point, you have several options. You can:
 - Enter the fully qualified domain name of a Mailbox server in the source forest that can act as the remote MRS proxy server and then click **Next** to have Exchange try to connect to a migration endpoint on this server and then test connectivity.
 - Click **Back** to provide alternate credentials and then click **Next** to retry the connection with those credentials. (Or simply click **Back** and then click **Next** to retry the connection with the original credentials.)

- Use the Exchange Remote Connectivity Analyzer (<https://testexchangeconnectivity.com>) to diagnose the connectivity issues. Once the issues are resolved, you can click Next to continue.
7. On the Start The Batch page, your current login is selected as the recipient for the batch report. This report will contain details about errors encountered during the migration. To add or change recipients for this report, click **Browse**. Then use the Select Members dialog box to select the recipients that should receive the report and then click **OK**. You must select at least one recipient.
 8. By default, Exchange Server creates and starts the batch migration request. When the request is completed, Exchange Server will also automatically finalize it. If you want to manually start the batch, select the Manual option. If you want to manually finalize the batch, clear the Automatically Complete check box.
 9. Click **New**. Migrating mailboxes can take several hours, depending on the size and number of the mailboxes you are moving. You can check the status of move requests by refreshing the view on the Migration page. While the request is in the Synced state, you can cancel the request by selecting it and then clicking Delete. You cannot cancel a request that has started syncing.

You can perform online moves in Exchange Management Shell by using `New-MoveRequest` for each mailbox you plan to move. The following examples show two different ways you can move Adam Carpenter's mailbox:

```
New-MoveRequest -Identity 'adamc' -Remote
-RemoteHost 'mailserver62.imaginedlands.com'-mrserver
'mailserver19.imaginedlands.com'
-TargetDatabase "Engineering Primary"
```

```
'imaginedlands.com/users/Adam Carpenter' | New-MoveRequest -Remote
-RemoteHost 'mailserver62.imaginedlands.com' -mrserver
'mailserver19.imaginedlands.com'
-TargetDatabase 'Engineering Primary'
```

After you initiate a move, you can check the status of the online move by using `Get-MoveRequest`. As shown in the following example, the key parameters to provide are the identity of the mailbox you want to check and the name of the proxy server:

```
Get-MoveRequest -Identity 'adamc'
-mrserver 'mailserver19.imaginedlands.com'
```

You can use `Suspend-MoveRequest` to suspend a move request that is not yet complete, and `Resume-MoveRequest` to resume a suspended move request. Resuming a suspended request allows it to complete.

At any time prior to running the move request completing, you can cancel the move by running `Remove-MoveRequest` and specifying the identity of the mailbox that shouldn't be moved, such as:

```
Remove-MoveRequest -Identity 'adamc' -mrserver
```


'mailserver19.imaginedlands.com'

Managing Delivery Restrictions, Permissions, and Storage Limits

You use mailbox properties to set delivery restrictions, permissions, and storage limits. To change these configuration settings for mailboxes, follow the techniques discussed in this section.

Setting Message Size Restrictions for Contacts

You set message size restrictions for contacts in much the same way that you set size restrictions for users. Follow the steps listed in the next section.

Setting Message Size Restrictions on Delivery to and from Individual Mailboxes

Message size restrictions control the maximum size of messages that can be sent or received in the Exchange organization. With Exchange Online, the maximum size of messages that users can send is 35,840 KB and the maximum size of messages that users can receive is 36,864 KB by default. With on-premises Exchange, you can manage these settings in a variety of ways. Typically, you manage these restrictions for the organization as a whole using the Organization Transport Settings. To manage these settings complete these steps:

1. In Exchange Admin Center, select **Mail Flow** in the Navigation menu and then select **Receive Connectors**.
2. On the Receive Connectors page, click More (**⋮**) and then select **Organization Transport Settings**.
3. By default, the maximum receive and send message size are both set to 10 MB. Use the options on the Limits page to set new defaults and then click **Save**.



The screenshot shows the 'organization transport settings' page with a left-hand navigation menu containing 'limits', 'safety net', and 'delivery'. The 'limits' section is active and contains three dropdown menus: 'Maximum number of recipients' set to 500, 'Maximum receive message size (MB)' set to 10, and 'Maximum send message size (MB)' set to 10. At the bottom right, there are 'Save' and 'Cancel' buttons.

You also can manage these restrictions using transport rules that filter messages by size and have specific conditions that apply to the size of messages or attachments, including the Apply This Rule If The Message Size Is Greater Than Or Equal To condition and the

Apply This Rule If Any Attachment Is Greater Than Or Equal To condition.

new rule

Name:
Message Size Filter

*Apply this rule if...
The message size is greater than or equal to... 40.00 MB
add condition

*Do the following...
Reject the message with the explanation... 'Message is too large, at 40 MB or larger.'
add action

Except if...
add exception

Properties of this rule:
 Audit this rule with severity level:
Low

Choose a mode for this rule:
 Enforce

Save Cancel

Using the Apply This Rule If The Message Size Is Greater Than Or Equal To condition, you can:

- Set restrictions regarding the size of messages that can be sent or received.
- Specify the action or actions to take if a message meets or exceeds this limit.
- Define exceptions for specific users and groups as well as for messages that have specifically-defined characteristics.

In Exchange Admin Center, you can create and manage transport rules, using the options found under Mail Flow > Rules. Click **New** and then select **Filter Messages By Size**.

The shell commands for working with transport rules include: Disable-TransportRule, Enable-TransportRule, Get-TransportRule, New-TransportRule, Remove-TransportRule, and Set-TransportRule.

When setting these types of organization-wide restrictions, you'll want to consider the global impact. Typically, you'll want to apply organization-wide restrictions only to prevent abuse of the mail system. For example, you may want to configure rules that block sending and receiving of very large files and provide a message that encourages senders to use a site mailbox configured as part of a Microsoft SharePoint site for sharing large documents instead.

Sometimes, you need to set exceptions for specific users. For example, some users might need to be able to send large files as part of their job.

While no delivery restrictions are set by default with on-premises Exchange, specific restrictions are set for online Exchange by default. For sending messages, the maximum message size is 35840 KB. For received messages, the maximum message size is 36864 KB. You can override these defaults by setting different maximum send and receive

sizes, up to 153600 KB.

You set individual delivery restrictions by completing the following steps:

1. Open the Properties dialog box for the mailbox-enabled user account by double-clicking the user name in Exchange Admin Center under Recipients > Mailboxes.
2. On the Mailbox Features page, scroll down and then click **View Details** under Message Size Restrictions.
3. As shown in Figure 7-11, you can set the following send and receive restrictions:
 - **Sent Messages > Maximum Message Size** Sets a limit on the size of messages the user can send. The value is set in kilobytes (KBs). If an outgoing message exceeds the limit, the message isn't sent and the user receives a non-delivery report (NDR).
 - **Received Messages > Maximum Message Size** Sets a limit on the size of messages the user can receive. The value is set in KBs. If an incoming message exceeds the limit, the message isn't delivered and the sender receives an NDR.

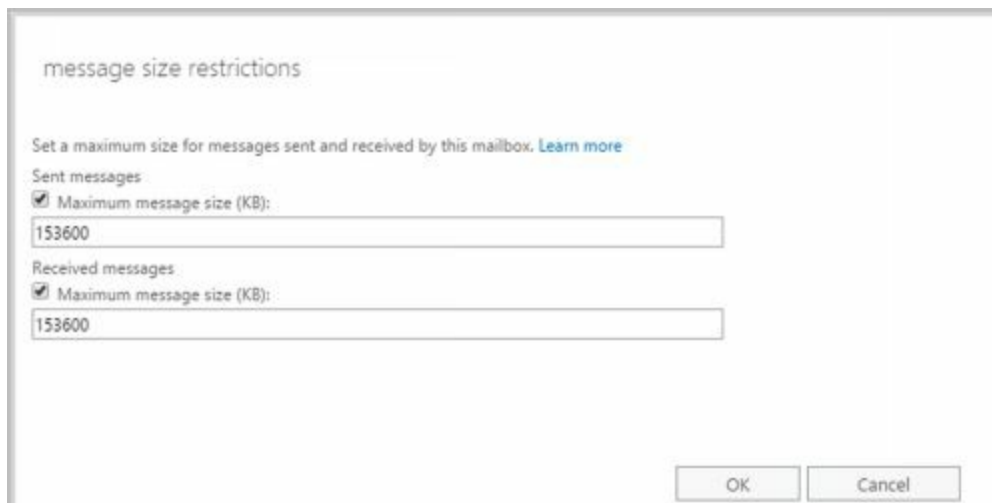


FIGURE 7-11 You can apply individual delivery restrictions on a per-user basis.

4. Click **OK** and then click **Save**. The restrictions that you set override the global default settings.

Setting Send and Receive Restrictions for Contacts

You set message send and receive restrictions for contacts in the same way that you set these restrictions for users. Follow the steps listed in the next section.

Setting Message Send and Receive Restrictions on Individual Mailboxes

By default, user mailboxes are configured to accept messages from anyone. To override this behavior, you can do the following:

- Specify that only messages from the listed users, contacts, or groups be accepted.
- Specify that messages from specific users, contacts, or groups be rejected.

- Specify that only messages from authenticated users—meaning users who have logged on to the Exchange system or the domain—be accepted.

With both on-premises Exchange and Exchange Online, you set message send and receive restrictions by completing the following steps:

1. Open the Properties dialog box for the mailbox-enabled user account by double-clicking the user name in Exchange Admin Center under Recipients > Mailboxes.
2. On the Mailbox Features page, scroll down and then click **View Details** under Message Delivery Restrictions. As shown in Figure 7-12, you can then set message acceptance restrictions.

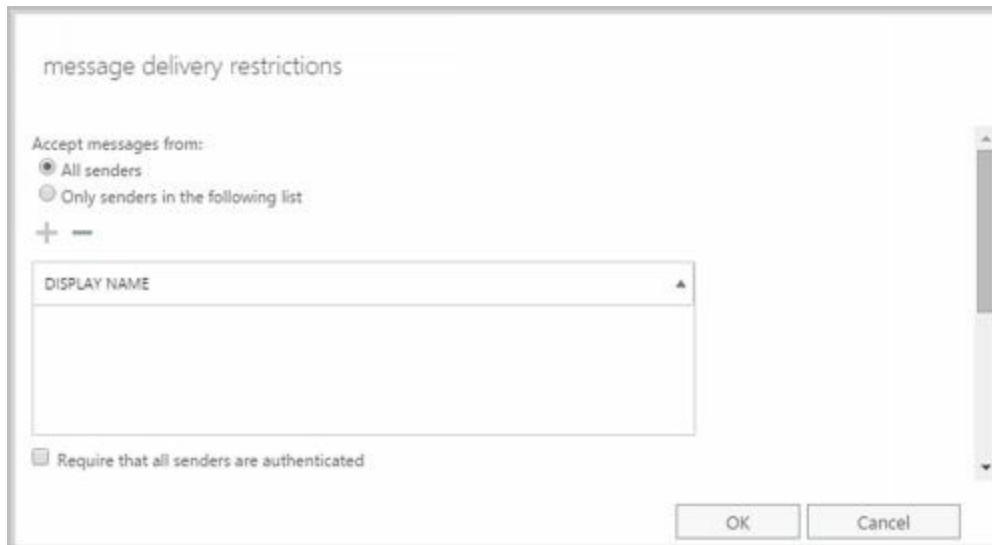



FIGURE 7-12 You can apply send and receive restrictions on messages on a per-user basis.

3. To accept messages from all email addresses except those on the reject list, under Accept Messages From, select **All Senders**.
4. To specify that only messages from the listed users, contacts, or groups be accepted, select the **Only Senders In The Following List** option and then add acceptable recipients by following these steps:

- Click Add () to display the Select Members dialog box.
- Select a recipient, and then click OK. Repeat as necessary.

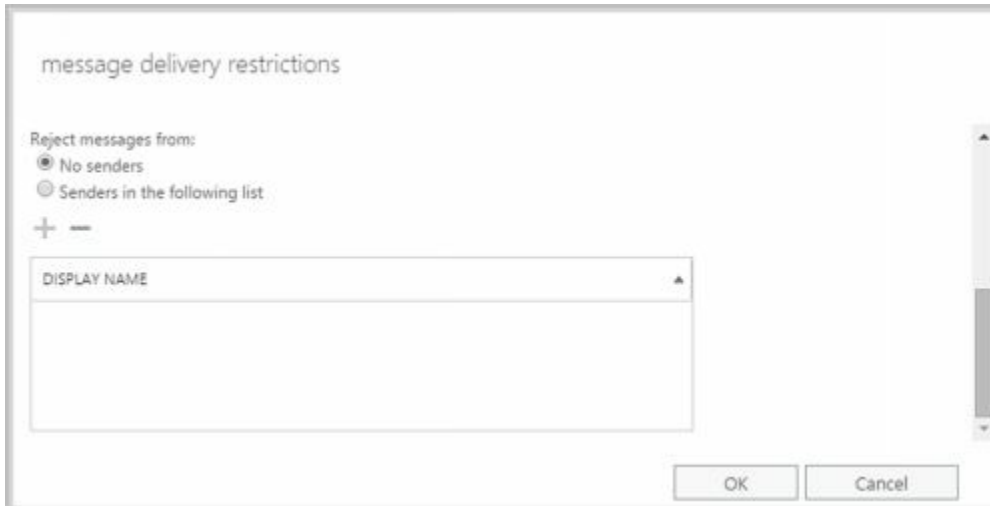
TIP You can select multiple recipients at the same time. To select multiple recipients individually, hold down the Ctrl key and then click each recipient that you want to select. To select a sequence of recipients, select the first recipient, hold down the Shift key, and then click the last recipient.

5. If you want to ensure that messages are accepted only from authenticated users, select the **Require That All Senders Are Authenticated** check box.
6. To specify that no recipients should be rejected, under Reject Messages From, select **No Senders**.
7. To reject messages from specific recipients, under Reject Messages From, select

Senders In The Following List and then add unacceptable recipients by following these steps:

- Click Add () to display the Select Members dialog box.
- Select a recipient, and then click OK. Repeat as necessary

8. Click **OK**.




Permitting Others to Access a Mailbox

Occasionally, users need to access someone else's mailbox, and in certain situations, you should allow this. For example, if John is Susan's manager and Susan is going on vacation, John might need access to her mailbox while she's away. Another situation in which someone might need access to another mailbox is when you've set up special-purpose mailboxes, such as a mailbox for `Webmaster@domain.com` or a mailbox for `Info@domain.com`.


You can grant permissions for a mailbox in three ways:

- You can grant access to a mailbox and its content. If you want to grant access to a mailbox and its contents but not grant Send As permissions, use the Full Access settings. In Exchange Admin Center, open the Properties dialog box for the mailbox you want to work with and then select Mailbox Delegation. On the Mailbox

Delegation page, under Full Access, click Add (), and then use the Select Full Access dialog box to choose the recipients who should have access to the mailbox. To revoke the authority to access the mailbox, select an existing user name in the Display Name list box and then click Remove.




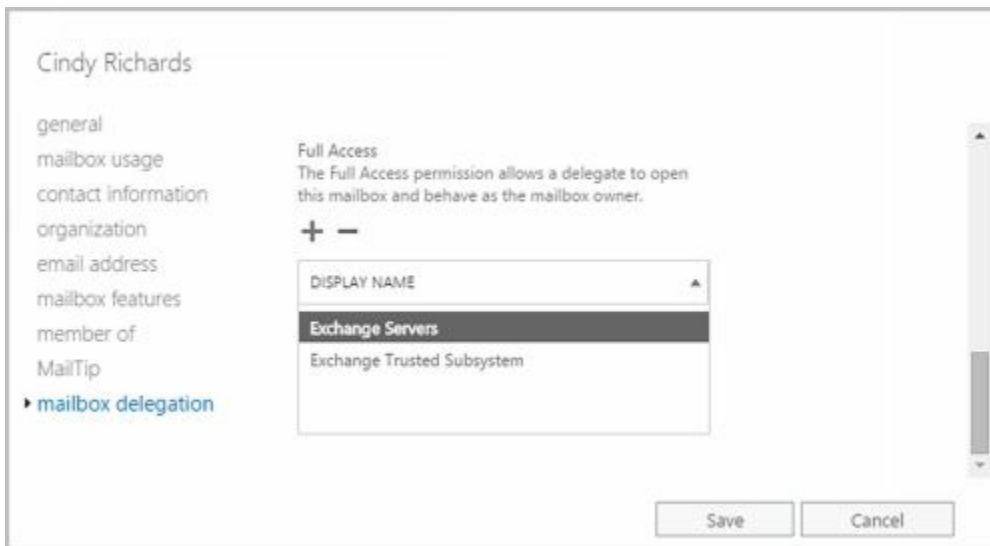
- You can grant the right to send messages as the mailbox owner. If you want to grant Send As permissions, use the Send As settings. In Exchange Admin Center, open the Properties dialog box for the mailbox you want to work with and then select Mailbox

Delegation. On the Mailbox Delegation page, under Send As, click Add (), and then use the Select Send As dialog box to choose the recipients who should have this permission. To revoke this permission, select an existing user name in the Display Name list box and then click Remove.



- You can grant the right to send messages on behalf of the mailbox owner. If you want to allow a user to send messages from a user's mailbox but want recipients to know a message was sent on behalf of the mailbox owner (rather than by the mailbox owner), grant Send On Behalf Of permissions. In Exchange Admin Center, open the Properties dialog box for the mailbox, and then select Mailbox Delegation. On the Mailbox Delegation page, under Send On Behalf Of, click Add

(), and then use the Select Send On Behalf Of dialog box to choose the recipients who should have this permission. To revoke this permission, select an existing user name in the Display Name list box and then click Remove.



In Exchange Management Shell, you can use the `Add-MailboxPermission` and `Remove-MailboxPermission` cmdlets to manage full access permissions. Samples 7-2 and 7-3 show examples of using these cmdlets. In these examples, the `AccessRights` parameter is set to `FullAccess` to indicate full access permissions on the mailbox.

SAMPLE 7-2 Adding full access permissions

Syntax

```
Add-MailboxPermission -Identity UserBeingGrantedPermission
-User UserWhoseMailboxIsBeingConfigured -AccessRights 'FullAccess'
```

Usage

```
Add-MailboxPermission -Identity
' CN=Mike Lam,OU=Engineering,DC=pocket-consultant,DC=com '
-User 'IMAGINEDLANDS\boba' -AccessRights ' FullAccess '
```

SAMPLE 7-3 Removing full access permissions

Syntax

```
Remove-MailboxPermission -Identity ' UserBeingGrantedPermission'
-User 'UserWhose MailboxIsBeingConfigured' -AccessRights 'FullAccess'
-InheritanceType 'All'
```

Usage

```
Remove-MailboxPermission -Identity ' CN=Jerry Orman,
OU=Engineering,DC=pocket-consultant,DC=com '
-User 'IMAGINEDLANDS\boba' -AccessRights ' FullAccess ' -InheritanceType 'All'
```

In Exchange Management Shell, you can use the `Add-ADPermission` and `Remove-ADPermission` cmdlets to manage Send As permissions. Samples 7-4 and 7-5 show examples using these cmdlets. In these examples, the `-ExtendedRights` parameter is set to `Send-As` to indicate you are setting Send As permissions on the mailbox.

SAMPLE 7-4 Adding send as permissions

Syntax

Add-ADPermission –Identity **UserBeingGrantedPermission**
–User **UserWhoseMailboxIsBeingConfigured** –ExtendedRights 'Send-As'

Usage

Add-ADPermission –Identity 'CN=Jerry
Orman,OU=Engineering,DC=cpandl,DC=com '
–User '**IMAGINEDLANDS\boba**' –ExtendedRights ' Send-As '

SAMPLE 7-5 Removing send as permissions

Syntax

Remove-ADPermission –Identity **UserBeingRevokedPermission**
–User **UserWhoseMailboxIsBeingConfigured** –ExtendedRights 'Send-As'
–InheritanceType 'All' –ChildObjectTypes \$null
–InheritedObjectType \$null -Properties \$null

Usage

Remove-ADPermission –Identity 'CN=Jerry
Orman,OU=Engineering, DC=pocket-consultant,DC=com'
–User '**IMAGINEDLANDS\boba**' –ExtendedRights 'Send-As'
–InheritanceType 'All' –ChildObjectTypes \$null –InheritedObjectTypes \$null
-Properties \$null

NOTE Another way to grant access permissions to mailboxes is to do so through Outlook. Using Outlook, you have more granular control over permissions. You can allow a user to log on as the mailbox owner, delegate mailbox access, and grant various levels of access. For more information on this issue, see the “Accessing Multiple Exchange Mailboxes” and “Granting Permission to Access Folders Without Delegating Access” sections in Chapter 10 “Configuring Exchange Clients.”

Forwarding Email to a New Address

Except when rights management prevents it, any messages sent to a user’s mailbox can be forwarded to another recipient. This recipient can be another user or a mail-enabled contact. To configure mail forwarding, follow these steps:

1. Open the Properties dialog box for the mailbox-enabled user account by double-clicking the user name in Exchange Admin Center.
2. On the Mailbox Features page, scroll down and then click **View Details** under Mail Flow.

delivery options

Forwarding Address
Forward email to the following recipient. [Learn more](#)

Enable forwarding
Forward email to the following recipient:

Mark Vance X Browse...

Deliver message to both forwarding address and mailbox

Recipient limit
 Maximum recipients:

OK Cancel

3. To remove forwarding, clear the Enable Forwarding check box.
4. To add forwarding, select the Enable Forwarding check box and then click **Browse**. Use the Select Mailbox User And Mailbox dialog box to choose the alternate recipient.
5. If you enabled forward, you can optionally specify that copies of forwarded messages should be retained in the original mailbox by selecting the **Deliver Message To Both Forwarding Address And Mailbox** checkbox.

If you use Exchange Management Shell to configure forwarding, you can specify that messages should be delivered to both the forwarding address and the current mailbox by setting the `-DeliverToMailboxAndForward` parameter to `$true` when using `Set-Mailbox`.

Setting Storage Restrictions on Mailbox and Archives

In a standard configuration of Exchange Online, each licensed user gets 25 GB of mailbox storage and a storage warning is issued when the mailbox reaches 22.5 GB. Similarly, if user has a licensed in-place archive, the archive can have up to 25 GB of storage; a storage warning is issued when the archive mailbox reaches 22.5 GB. Other licensing options are available that may grant additional storage rights.

With on-premises Exchange, you can set storage restrictions on multiple mailboxes using global settings for each mailbox database or on individual mailboxes using per-user restrictions. Global restrictions are applied when you create a mailbox and are reapplied when you define new global storage restrictions. Per-user storage restrictions are set individually for each mailbox and override the global default settings. By default, users can store up to 2 GB in their mailboxes. The quotas are set to:

- Issue a warning when the mailbox reaches 1.9 GB
- Prohibit send when the mailbox reaches 2 GB
- Prohibit send and receive when the mailbox reaches 2.3 GB

In contrast, the default settings for archive mailboxes allow users to store up to 50 GB in their archive mailboxes, and a warning is issued when the archive mailbox reaches 45 GB.

NOTE Storage restrictions apply only to mailboxes stored on the server. They don't apply to personal folders. Personal folders are stored on the user's computer

To configure global storage restrictions, you edit the properties of mailbox databases. In Exchange Admin Center, navigate to Servers > Databases. Open the Properties dialog box for the mailbox database by double-clicking the database name. On the Limits page, set the desired storage restrictions using the options provided.



You set individual storage restrictions for mailboxes by completing the following steps:

1. Open the Properties dialog box for the mailbox-enabled user account by double-clicking the user name in Exchange Admin Center.
2. On the Mailbox Usage page, click **More Options**. You'll then see the storage restrictions as shown in Figure 7-13.

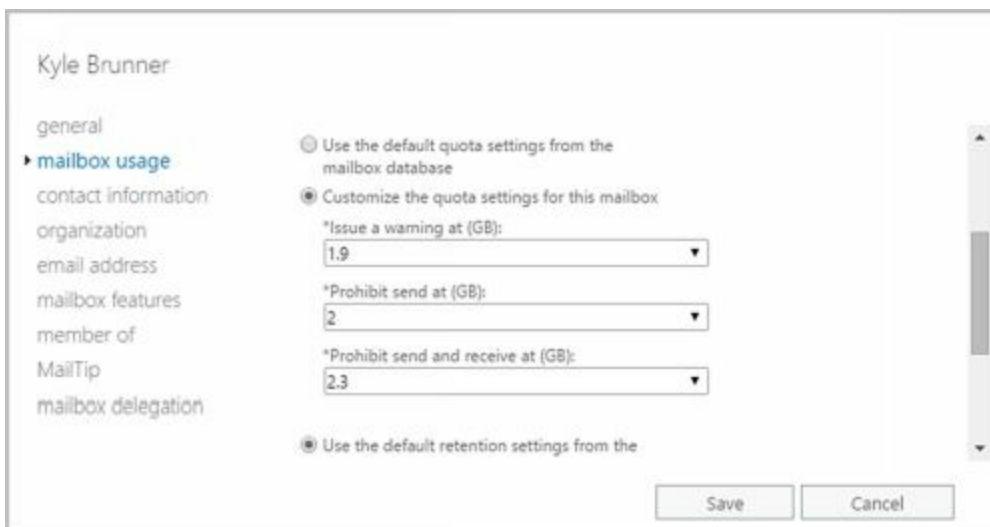


FIGURE 7-13 Use the quota settings to specify storage limits and deleted item retention on a per-user basis when necessary.

3. To set mailbox storage limits, select **Customize The Quota Settings For This Mailbox**. Then set one or more of the following storage limits:
 - **Issue Warning At (GB)** This limit specifies the size, in gigabytes, that a mailbox can reach before a warning is issued to the user. The warning tells the user to clean out the mailbox.
 - **Prohibit Send At (GB)** This limit specifies the size, 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 mailbox size is under the limit.

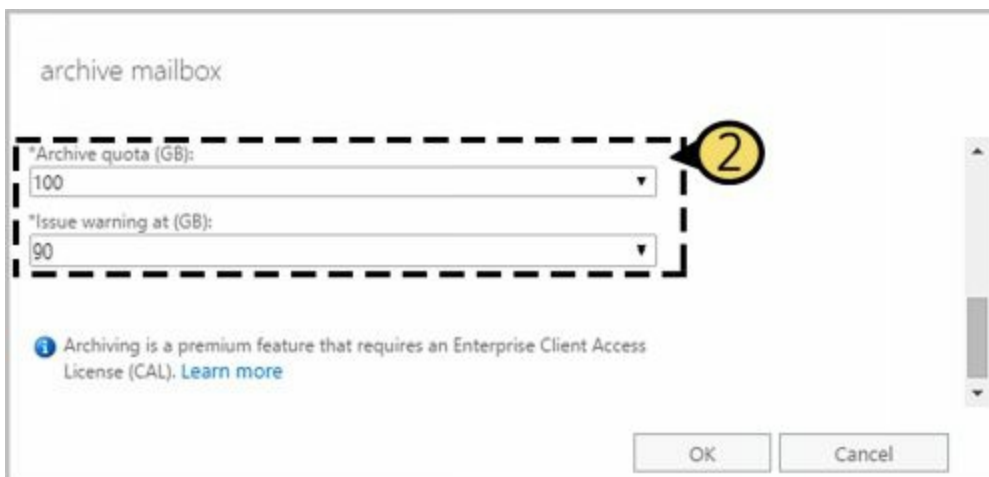
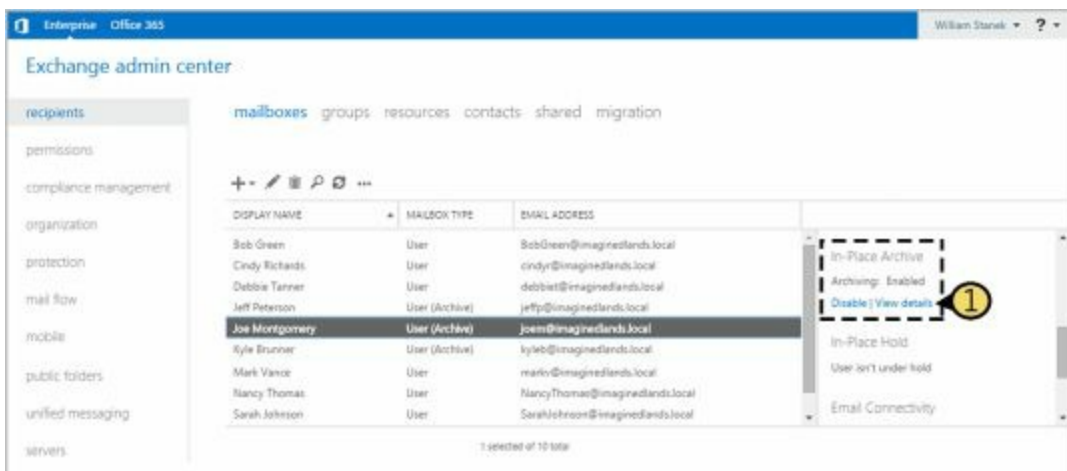
- **Prohibit Send And Receive At (GB)** This limit specifies the size, in gigabytes, that a mailbox can reach before the user is prohibited from sending and receiving mail. The restriction ends when the user clears out the mailbox and the mailbox size is under the limit.

CAUTION Prohibiting send and receive might cause the user to think they've lost email. When someone sends a message to a user who is prohibited from receiving messages, an NDR is generated and delivered to the sender. The original recipient never sees the email. Because of this, you should rarely prohibit send and receive.

4. Click Save.

Users who have an archive mailbox have the mailbox type User (Archive). You set individual storage restrictions for archive mailboxes by completing the following steps:

1. Select the user name in Exchange Admin Center.
2. In the Details pane, scroll down until you see the In-Place Archive heading and the related options. Click **View Details**.
3. Enter the desired maximum size for the archive in the Archive Quota text box.
4. Enter the storage limit for issuing a storage warning in the Issue Warning At text box.
5. Click **OK**.



Setting Deleted Item Retention Time on Individual Mailboxes

@techjob

Normally, when a user deletes a message in Outlook, the message is placed in the Deleted Items folder. The message remains in the Deleted Items folder until the user deletes it manually or allows Outlook to clear out the Deleted Items folder. With personal folders, the message is then permanently deleted and you can't restore it. With server-based mailboxes, the message isn't actually deleted from the Exchange database. Instead, the message is marked as hidden and kept for a specified period of time called the *deleted item retention period*.

NOTE The standard processes can be modified in several different ways. A user could press Shift+Delete to bypass Deleted Items. As an administrator, you can create and apply policies that prevent users from deleting items (even if they try to use Shift+Delete). You can also configure policy to retain items indefinitely.

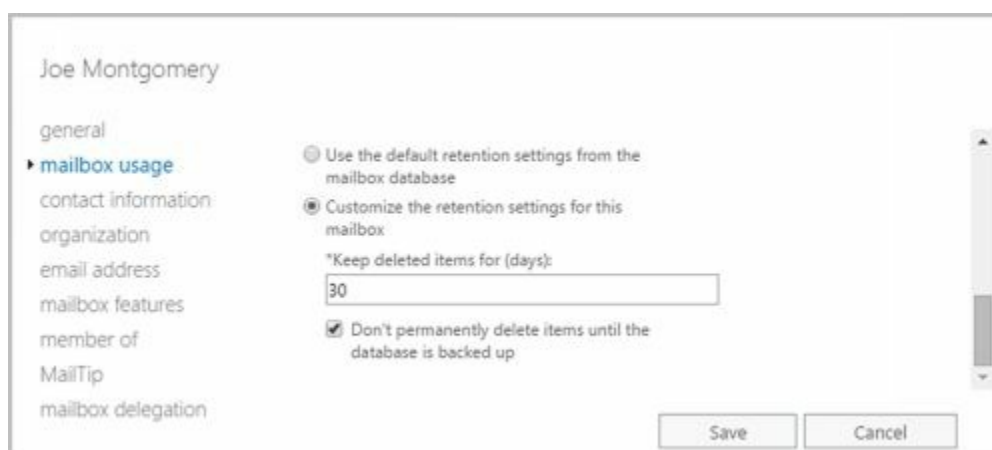
Default retention settings are configured for each mailbox database in the organization. With Exchange Online, the retention settings are as follows:

- Deleted items are retained for a maximum of 30 days.
- Items removed from the Deleted Items folder are retained for a maximum of 14 days.
- Items in the Junk Folder are retained for a maximum of 30 days before they are removed.

To configure deleted item retention on a per database basis, you edit the properties of mailbox databases. In Exchange Admin Center, navigate to Servers > Databases. Open the Properties dialog box for the mailbox database by double-clicking the database name. On the Limits page, use the options provided to configure the deleted item retention settings.

You can override the database settings on a per-user basis by completing these steps:

1. Open the Properties dialog box for the mailbox-enabled user account by double-clicking the user name in Exchange Admin Center.
2. On the Mailbox Usage page, click **More Options** and then select **Customize The Retention Settings For This Mailbox**.



The screenshot shows the 'Mailbox Usage' settings for a user named 'Joe Montgomery'. The 'mailbox usage' section is expanded, and the 'Limits' sub-section is visible. There are two radio button options: 'Use the default retention settings from the mailbox database' (which is unselected) and 'Customize the retention settings for this mailbox' (which is selected). Below these options, there is a text input field labeled '*Keep deleted items for (days):' with the value '30' entered. At the bottom of this section, there is a checked checkbox labeled 'Don't permanently delete items until the database is backed up'. At the very bottom of the dialog box, there are 'Save' and 'Cancel' buttons.

3. In the Keep Deleted Items For (Days) text box, enter the number of days to retain deleted items. An average retention period is 14 days. If you set the retention period to 0 and aren't using policies that prevent deletion, messages aren't retained and can't be recovered. If you set the retention period to 0 but are using policies that prevent deletion, the messages are retained according to the established policies.
4. You can also specify that deleted messages should not be permanently removed until the mailbox database has been backed up. This option ensures that the deleted items are archived into at least one backup set. Click **Save**.

REAL WORLD Deleted item retention is convenient because it allows the administrator the chance to salvage accidentally deleted email without restoring a user's mailbox from backup. I strongly recommend that you enable this setting, either in the mailbox database or for individual mailboxes, and configure the retention period accordingly.

Chapter 8. Managing Groups

Learning the ins and outs of distribution groups and address lists will greatly facilitate the efficiency and effectiveness of Microsoft Exchange Server and Exchange Online administration. Careful planning of your organization's groups and address lists can save you time and help your organization run more efficiently. Study the concepts discussed in this chapter and the next then use the step-by-step procedures to implement the groups and lists for your organization.

Using Security and Distribution Groups

You use groups to grant permissions to similar types of users, to simplify account administration, and to make it easier to contact multiple users. For example, you can send a message addressed to a group, and the message will go to all the users in that group. Thus, instead of having to enter 20 different email addresses in the message header, you enter one email address for all of the group members.

Group Types, Scope, And Identifiers

Windows defines several different types of groups, and each of these groups can have a unique scope. In Active Directory domains, you use three group types:

- **Security** You use security groups to control access to network resources. You can also use user-defined security groups to distribute email.
- **Standard distribution** Standard distribution groups have fixed membership, and you use them only as email distribution lists. You can't use these groups to control access to network resources.
- **Dynamic distribution** Membership for dynamic distribution groups is determined based on a Lightweight Directory Access Protocol (LDAP) query; you use these groups only as email distribution lists. The LDAP query is used to build the list of members whenever messages are sent to the group.

Security groups can have different scopes—*domain local*, *global*, and *universal*—so that they are valid in different areas of your Active Directory forest. Exchange Server only supports groups with universal scope. You can mail-enable security groups with universal scope, and you can create new distribution groups with universal scope.

In Exchange Admin Center, you select Recipients in the Navigation menu and then select Groups to work with groups (see Figure 8-1). Only mail-enabled groups with universal scope are displayed. Groups with universal scope can do the following:

- Contain users and groups from any domain in the Active Directory forest
- Be added to other groups and assigned permissions in any domain in the Active Directory forest

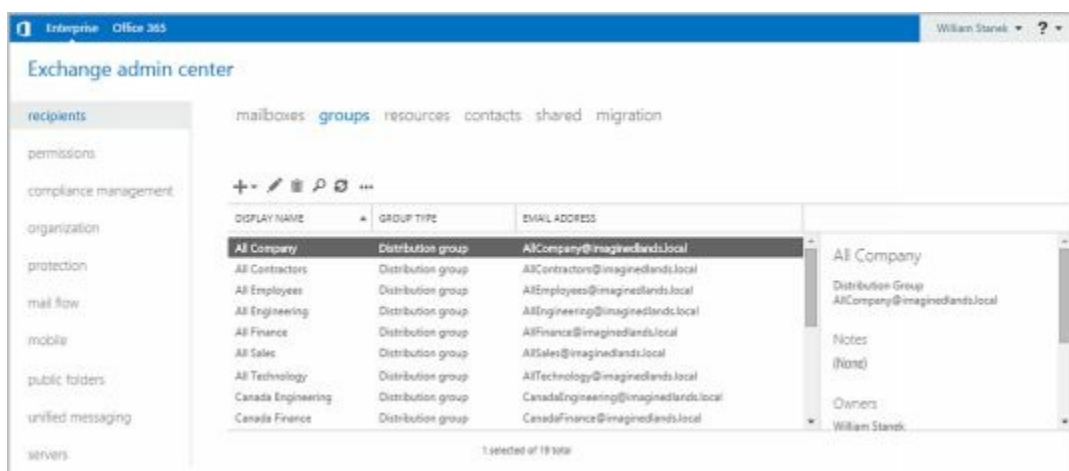


FIGURE 8-1 Viewing the configured groups in Exchange Admin Center.

When you work with dynamic distribution groups, keep in mind that the membership can include only members of the local domain, or it can include users and groups from other domains, domain trees, or forests. Scope is determined by the default apply-filter container you associate with the group when you create it. More specifically, the default apply-filter container defines the root of the search hierarchy and the LDAP query filters to recipients in and below the specified container. For example, if the apply-filter container you associate with the group is `imaginedlands.com`, the query filter is applied to all recipients in this domain. If the apply-filter container you associate with the organizational unit is `Engineering`, the query filter is applied to all recipients in or below this container.

As with user accounts, Windows uses unique security identifiers (SIDs) to track groups. This means that you can't delete a group, re-create it with the same name, and then expect all the permissions and privileges to remain the same. The new group will have a new SID, and all the permissions and privileges of the old group will be lost.

When to Use Security and Standard Distribution Groups

Rather than duplicating your existing security group structure with distribution groups that have the same purpose, you might want to selectively mail-enable your universal security groups, which converts them to distribution groups. For example, if you have a universal security group called `Marketing`, you don't need to create a `MarketingDistList` distribution group. Instead, you could enable Exchange mail on the original universal security group, which would then become a distribution group.

You might also want to mail-enable universal security groups that you previously defined. Then, if existing distribution groups serve the same purpose, you can delete the distribution groups.

To reduce the time administrators spend managing groups, Exchange defines several additional control settings, including

- **Group ownership** Mail-enabled security groups, standard distribution groups, and dynamic distribution groups can have one or more owners. A group's owners are the users assigned as its managers, and they can control membership in the group. A group's managers are listed when users view the properties of the group in Microsoft Office Outlook. Additionally, managers can receive delivery reports for groups if you select the `Send Delivery Reports To Group Manager` option when configuring group settings.
- **Membership approval** Mail-enabled security groups and standard distribution groups can have open or closed membership. There are separate settings for joining and leaving a group. For joining, the group can be open to allow users to join without requiring permission, be closed to allow only group owners and administrators to add members, or require owner approval to allow users to request membership in a group. Membership requests must be approved by a group owner. For leaving, a group can

either be open to allow users to leave a group without requiring owner approval or closed to allow only group owners and administrators to remove members.

Your management tool of choice will determine your options for configuring group ownership and membership approval. When you create groups in Exchange Admin Center, you can specify ownership, membership, and approval settings when you create the group and can edit these settings at any time by editing the group's properties. When you create groups in Exchange Management Shell, you can configure additional advanced options that you'd otherwise have to manage after creating the group in Exchange Admin Center.

When to Use Dynamic Distribution Groups

It's a fact of life that over time users will move to different departments, leave the company, or accept different responsibilities. With standard distribution groups, you'll spend a lot of time managing group membership when these types of changes occur—and that's where dynamic distribution groups come into the picture. With dynamic distribution groups, there isn't a fixed group membership and you don't have to add or remove users from groups. Instead, group membership is determined by the results of an LDAP query sent to your organization's Global Catalog.

Dynamic distribution groups can be used with or without a dedicated expansion server. You'll get the most benefit from dynamic distribution without a dedicated expansion server when the member list returned in the results is relatively small (fewer than 25 members). In the case of potentially hundreds or thousands of members, however, dynamic distribution is inefficient and could require a great deal of processing to complete. Exchange 2016 shifts the processing requirements from the Global Catalog server to a dedicated expansion server (a server whose only task is to expand the LDAP queries). By default, Exchange 2016 uses the closest Mailbox server as the dedicated expansion server. For more information on expansion servers, see "Designating an expansion server" later in this chapter.

One other thing to note about dynamic distribution is that you can associate only one specific query with each distribution group. For example, you could create separate groups for each department in the organization. You could have groups called QD-Accounting, QD-BizDev, QD-Engineering, QD-Marketing, QD-Operations, QD-Sales, and QD-Support. You could, in turn, create a standard distribution group or a dynamic distribution group called AllEmployees that contains these groups as members—thereby establishing a distribution group hierarchy.

When using multiple parameters with dynamic distribution, keep in mind that multiple parameters typically work as logical AND operations. For example, if you create a query with a parameter that matches all employees in the state of Washington with all employees in the Marketing department, the query results do not contain a list of all employees in Washington or all Marketing employees. Rather, the results contain a list of recipients who are in Washington and are members of the Marketing group. In this case, you get the expected results by creating a dynamic distribution group for all

Washington State employees, another dynamic distribution group for all Marketing employees, and a final group that has as members the other two distribution groups.

Working with Security and Standard Distribution Groups

As you set out to work with groups, you'll find that some tasks are specific to each type of group and some tasks can be performed with any type of group. Because of this, I've divided the group management discussion into three sections. In this section, you'll learn about the typical tasks you perform with security and standard distribution groups. The next section discusses tasks you'll perform only with dynamic distribution groups. The third section discusses general management tasks.

You can use Exchange Admin Center or Exchange Management Shell to work with groups.

Group Naming Policy

Whether you work at a small company with 50 employees or a large enterprise with 5,000 employees, you should consider establishing a group naming policy that ensures a consistent naming strategy is used for group names. For administrators, your naming policy should be implemented through written policies within your IT department and could be applied to both security groups and distribution groups.

Exchange 2016 and Exchange Online also allow you to establish official naming policy for standard distribution groups. Group naming policy is:

- Applied to non-administrators whenever they create or rename distribution groups.
- Applied to administrators only when they create or rename distribution groups using the shell (and omit the `-IgnoreNamingPolicy` parameter).

IMPORTANT Group naming policy doesn't apply to security groups or dynamic distribution groups. Each Exchange organization can have one and only one naming policy. Any naming policy you define is applied throughout the Exchange organization.

Understanding Group Naming Policy

You use group naming policy to format group names according to a defined standard. The rules for naming policy allow for one or more prefixes, a group name, and one or more suffixes, giving an expanded syntax of:

`<Prefix1><Prefix2>...<Prefix N><GroupName><Suffix1><Suffix2>...<Suffix N>`

You can use any Exchange attribute as the prefix or suffix. You also can use a text string as a prefix or suffix. The prefix, group name and suffix are combined without spacing. To improve readability, you can separate the prefix, name and suffix with a placeholder character, such as a space (), a period (.) or a dash (-).

Group naming policy works like this:

- A user creates a standard distribution group and specifies a display name for the group. After creating the group, Exchange applies the group naming policy by adding any prefixes or suffixes defined in the group naming policy to the display name.
- The display name is displayed in the distribution groups list in Exchange Admin Center, the shared address book, and the To:, Cc:, and From: fields in email messages.

You can create a naming policy with only a prefix and group name or with only a suffix and a group name. Common attributes that you might want to use as prefixes or suffixes include city, country code, department, office, and state. For example, you might want all distribution groups to have the following syntax:

State_ GroupName

To do this, you would create a naming policy with two prefixes. As shown in Figure 8-2, the first prefix would have the <State> attribute. The second prefix would have the _ text value. Thus, if a user in the state of New York (NY) creates a standard distribution group called Sales, Exchange adds the defined prefixes and the display name becomes NY_Sales.

FIGURE 8-2 Creating a naming policy with two prefixes.

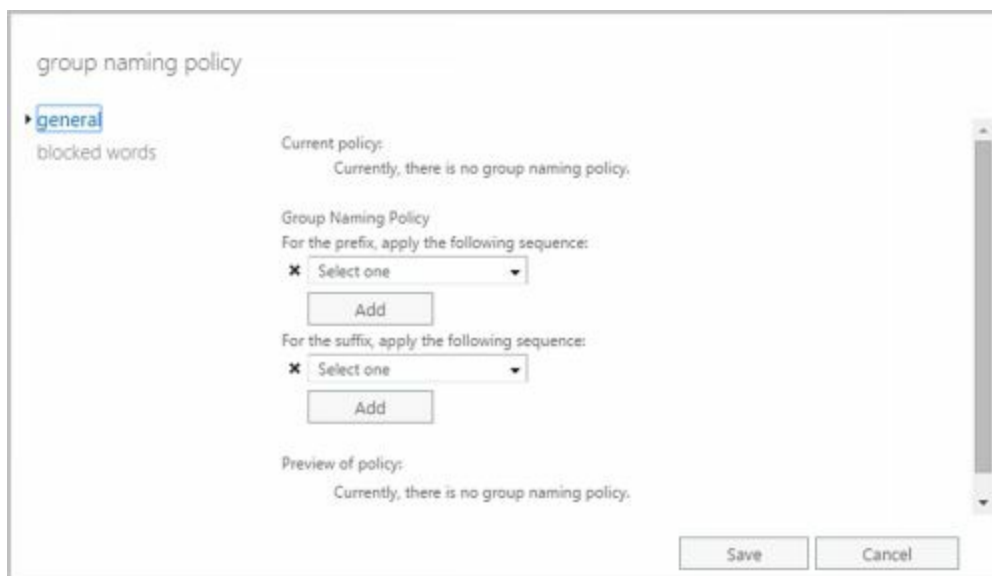
Group naming policy also allows you to specify blocked words. Users who try to use a word that you've blocked see an error message when they try to create the new group and are asked to remove the blocked word and create the group again.

Defining Group Naming Policy for Your Organization

Group naming policy formats display names so that they follow a defined standard. When setting the naming format, keep in mind that users enter the desired display name when they create the group and Exchange transforms the format according to the defined policy. Because the display name is limited to 64 characters, you must consider this limit when defining the prefixes and suffixes in your naming policy.

You can create the group naming policy for the Exchange organization by completing the following steps:

1. In Exchange Admin Center, select **Recipients** in the Navigation menu and then select **Groups** .
2. Click the **More** button (**⋮**) and then select **Configure Group Naming Policy** . This displays the Group Naming Policy dialog box.
3. If you want the naming policy to have a prefix, do one of the following and then optionally click **Add** to add additional prefixes using the same technique:
 - Use the selection list to choose **Attribute** as the prefix. In the **Select The Attribute** dialog box, select the attribute to use and then click **OK**.
 - Use the selection list to choose **Text** as the prefix. In the **Enter Text** dialog box, select the text string to use and then click **OK**.






4. If you want the naming policy to have a suffix do one of the following and then optionally click **Add** to add additional suffixes using the same technique:
 - Use the selection list to choose **Attribute** as the suffix. In the **Select The Attribute** dialog box, select the attribute to use and then click **OK**.
 - Use the selection list to choose **Text** as the suffix. In the **Enter Text** dialog box, select the text string to use and then click **OK**.
5. As you define the naming policy, the **Preview Of Policy** area shows the naming format. When you are satisfied with the naming format, click **Save** .

Defining Blocked Words in Group Naming Policy

Blocked words allow you to specify words that users can't use in the names of standard distribution groups they create. You can define or manage the blocked words list by completing the following steps:

1. In Exchange Admin Center, select **Recipients** in the Navigation menu and then select **Groups** .

2. Click the **More** button () and then select Configure Group Naming Policy. This displays the Group Naming Policy dialog box.
3. On the Blocked Words page, any currently blocked words are displayed. Use the following techniques to manage the blocked word list:
 - To add a blocked word, type the word in the text box provided and then click Add (). Alternatively, type the word to block in the text box provided and then press Enter.
 - To modify a blocked word, select the word in the blocked word list and then click Edit (). Modify the word and then click outside the text box provided for editing. Alternatively, press Enter to apply the edits.
 - To remove a blocked word, click the word to remove and then click Remove.
4. Click **Save** .



Creating Security and Standard Distribution Groups

Security groups and distribution groups are available whether you are working with online or on-premises Exchange organizations. You use groups to manage permissions and to distribute email. As you set out to create groups, remember that you create groups for similar types of users. Consequently, you might want to create the following types of groups:


- **Groups for departments within the organization** Generally, users who work in the same department need access to similar resources and should be a part of the same email distribution lists.
- **Groups for roles within the organization** You can also organize groups according to the users' roles within the organization. For example, you could use a group called Executives to send email to all the members of the executive team and a group called Managers to send email to all managers and executives in the organization.
- **Groups for users of specific projects** Often, users working on a major project need

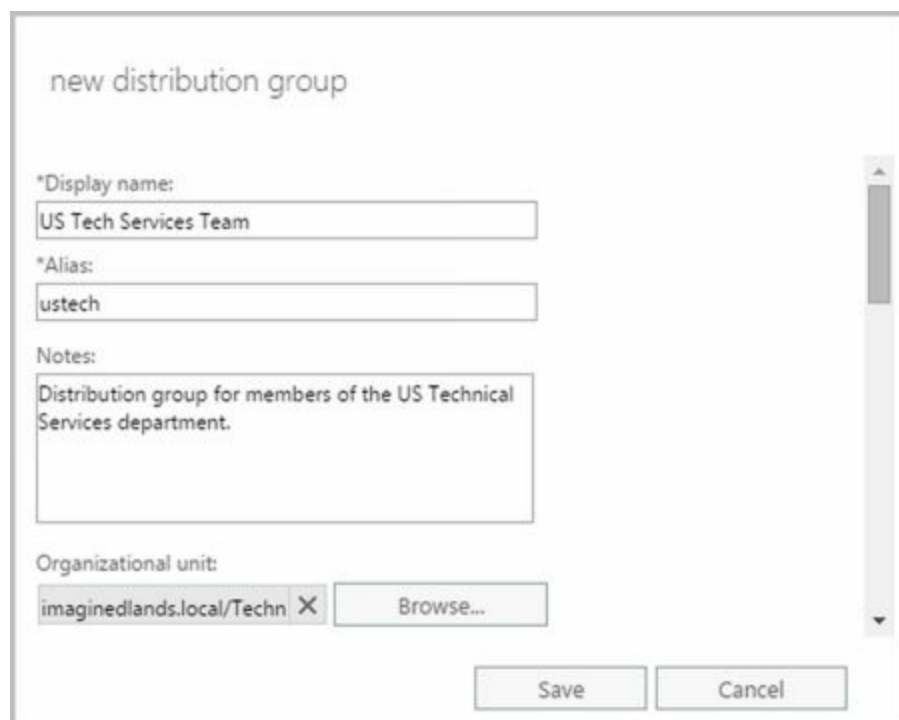
a way to send email to all the members of the team. To address this need, you can create a group specifically for the project.

You can create groups several ways. You can create a new distribution group, you can create a mail-enabled universal security group, or you can mail-enable an existing universal security group.

Creating a New Group

You can create a new distribution group or a new mail-enabled security group by completing the following steps:

1. In Exchange Admin Center, select **Recipients** in the Navigation menu and then select **Groups**.
2. Click New () and then do one of the following:
 - Select Distribution Group to create a new Distribution Group. This opens the New Distribution Group dialog box, shown in Figure 8-3.
 - Select Security Group to create a new mail-enabled Security Group. This opens the New Security Group dialog box, and the options are the same as those for new distribution groups.



The screenshot shows a dialog box titled "new distribution group". It contains the following fields and controls:

- *Display name:** A text box containing "US Tech Services Team".
- *Alias:** A text box containing "ustech".
- Notes:** A text area containing "Distribution group for members of the US Technical Services department."
- Organizational unit:** A text box containing "imaginedlands.local/Techn" with a close button (X) and a "Browse..." button.
- Buttons:** "Save" and "Cancel" buttons at the bottom.

FIGURE 8-3 Configuring the group's settings.

3. In the Display Name text box, type a display name for the group. Group names aren't case-sensitive and can be up to 64 characters long. Keep in mind that group naming policy doesn't apply to administrators creating distribution groups in Exchange Admin Center (or to mail-enabled security groups in any way).
4. Like users, groups have Exchange aliases. Enter an alias. The Exchange alias is used to set the group's SMTP email address. Exchange Server uses the SMTP

address for receiving messages.

5. For Exchange Online, the name and domain components of the default email address are displayed in the Email Address text boxes. As appropriate, change the default name and use the drop-down list to select the domain with which you want to associate the group. This sets the fully qualified email address, such as us-tech@imaginedlands.onmicrosoft.com.
6. With on-premises Exchange, the group account is created in the default user container, which typically is the Users container. To create the group in a specific organizational unit instead, click **Browse** to the right of the Organizational Unit text box. In the Select Organizational Unit dialog box, choose the location where you want to store the account and then click **OK**.
7. Group owners are responsible for managing a group. To add owners, under Owners, click **Add** (**+**). In the Select Owner dialog box, select users, groups, or both that should have management responsibility for the group. Select multiple users and groups using the Shift or Ctrl keys.


The screenshot shows a dialog box titled "new distribution group". It has two main sections: "Owners" and "Members".

- Owners:** Labeled "*Owners:", it contains a list with "William Stanek" selected. Above the list are "+" and "-" buttons.
- Members:** Labeled "Members:", it contains a checked checkbox "Add group owners as members" and a list with "US Technology" selected. Above the list are "+" and "-" buttons.

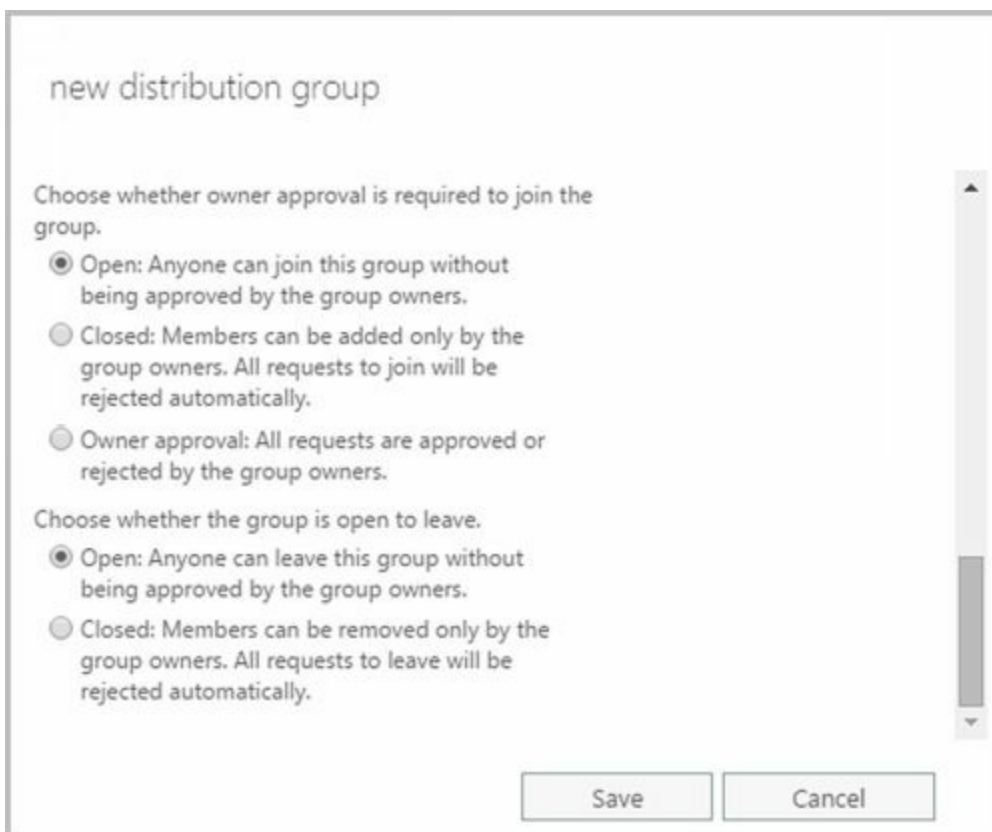
At the bottom of the dialog are "Save" and "Cancel" buttons.

IMPORTANT While dynamic distribution groups don't have to have owners, every mail-enabled security group and standard distribution group must have at least one owner. By default, the account you are using is set as the group owner.

8. Members of a group receive messages sent to the group. By default, the group owners are set as members of the group. If you don't want the currently listed owners to be members of the group, clear the **Add Group Owners As Members** checkbox.

9. To add members, under Members, click **Add** (). In the Select Members dialog box, select users, groups, or both that should be members of the group. Select multiple users and groups using the Shift or Ctrl keys.
10. Choose settings for joining the group. The options are:
- **Open** Anyone can join this group without being approved by the group owners.
 - **Closed** Members can be added only by the group owners. All requests to join will be rejected automatically.
 - **Owner Approval** All requests are approved or rejected by the group owners.
11. Choose settings for leaving the group. The options are:
- **Open** Anyone can leave this group without being approved by the group owners.
 - **Closed** Members can be removed only by the group owners. All requests to leave will be rejected automatically.
12. Click **Save** to create the group. If an error occurs during group creation, the related group will not be created. You need to correct the problem before you can complete this procedure. After creating a group, you might want to do the following:
- Set message size restrictions for messages mailed to the group.
 - Limit users who can send to the group.
 - Change or remove default email addresses.
 - Add more email addresses.

NOTE By default, the new distribution group is open for joining and open for leaving.



The screenshot shows a dialog box titled "new distribution group". It contains two sections of radio button options. The first section is titled "Choose whether owner approval is required to join the group." and has three options: "Open: Anyone can join this group without being approved by the group owners." (selected), "Closed: Members can be added only by the group owners. All requests to join will be rejected automatically.", and "Owner approval: All requests are approved or rejected by the group owners." The second section is titled "Choose whether the group is open to leave." and has two options: "Open: Anyone can leave this group without being approved by the group owners." (selected) and "Closed: Members can be removed only by the group owners. All requests to leave will be rejected automatically." At the bottom of the dialog are "Save" and "Cancel" buttons.

In Exchange Management Shell, you can create a new distribution group using the `New-DistributionGroup` cmdlet. Sample 8-1 provides the syntax and usage. You can set the `-Type` parameter to `Distribution` for a distribution group or to `Security` for a mail-enabled security group.

SAMPLE 8-1 `New-DistributionGroup` cmdlet syntax and usage

Syntax

```
New-DistributionGroup -Name ExchangeName [-Alias ExchangeAlias ]  
[-DisplayName DisplayName ] [-OrganizationalUnit OUName ]  
[-PrimarySmtpAddress SmtpAddress ] [-SamAccountName PreWin2000Name ]  
[-Type <Distribution | Security>] {AddtlParams}
```

```
{AddtlParams}  
[-ArbitrationMailbox ModeratorMailbox ] [-BypassNestedModerationEnabled  
<$true | $false>] [-CopyOwnerToMember {$true | $false}] [-DomainController  
FullyQualifiedName ] [-IgnoreNamingPolicy {$true | $false}] [-ManagedBy  
RecipientIdentities ] [-MemberDepartRestriction <Closed | Open |  
ApprovalRequired>] [-MemberJoinRestriction <Closed | Open |  
ApprovalRequired>] [-Members RecipientIdentities ] [-ModeratedBy  
Moderators ] [-ModerationEnabled <$true | $false>] [-Notes String ]  
[-Organization OrgName ] [-RoomList {$true | $false}]  
[-SendModerationNotifications <Never | Internal | Always>]
```

Usage

```
New-DistributionGroup -Name 'CorporateSales' -Type 'Distribution'  
-OrganizationalUnit 'imaginedlands.com/Sales'  
-SamAccountName 'CorporateSales'  
-DisplayName 'Corporate Sales'  
-Alias 'CorporateSales'
```

Mail-Enabling Universal Security Groups

You can't use Exchange Admin Center to mail-enable a security group. In Exchange Management Shell, you can mail-enable a universal security group using the `Enable-DistributionGroup` cmdlet. Sample 8-2 provides the syntax and usage.

SAMPLE 8-2 `Enable-DistributionGroup` cmdlet syntax and usage

Syntax

```
Enable-DistributionGroup -Identity GroupIdentity [-Alias ExchangeAlias ]  
[-DisplayName DisplayName ] [-DomainController FullyQualifiedName ]  
[-OverrideRecipientQuotas {$true | $false}]  
[-PrimarySmtpAddress SmtpAddress ]
```

Usage

```
Enable -DistributionGroup -Identity 'AllSales'  
-DisplayName 'All Sales' -Alias 'AllSales'
```

NOTE Group naming policy applies only to distribution groups.

You can manage mail-enabled security groups in several ways. You can add or remove group members as discussed in the “Assigning and Removing Membership for Individual Users, Groups, and Contacts” section of this chapter. If a group should no longer be mail-enabled, you can use `Disable-DistributionGroup` to remove the Exchange settings from the group. If you no longer need a mail-enabled security group and it is not a built-in group, you can permanently remove it from Active Directory by selecting it in Exchange Admin Center and clicking Delete. Alternatively, you can delete a group using `Delete-DistributionGroup`.

Using Exchange Management Shell, you can disable a group’s Exchange features using the `Disable-DistributionGroup` cmdlet, as shown in Sample 8-3.

SAMPLE 8-3 `Disable-DistributionGroup` cmdlet syntax and usage

Syntax

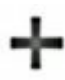

```
Disable-DistributionGroup -Identity GroupIdentity  
[-DomainController FullyQualifiedName ]  
[-IgnoreDefaultScope {$true | $false}]
```

Usage

```
Disable-DistributionGroup -Identity 'AllSales'
```

Assigning and Removing Membership for Individual Users, Groups, and Contacts

All users, groups, and contacts can be members of other groups. To configure a group’s membership, follow these steps:

1. In Exchange Admin Center, double-click the group entry. This opens the group’s Properties dialog box.
2. On the Membership page, you’ll see a list of current members. Click **Add** () to add recipients to the group. In the Select Members dialog box, select users, groups, or both that should be members of the group. Select multiple users and groups using the Shift or Ctrl keys.
3. You can remove members on the Membership page as well. To remove a member from a group, select a recipient, and then click **Remove** (). When you’re finished, click **Save** .



In Exchange Management Shell, you can view group members using the `Get-DistributionGroupMember` cmdlet. Sample 8-4 provides the syntax and usage.

SAMPLE 8-4 `Get-DistributionGroupMember` cmdlet syntax and usage

Syntax

```
Get-DistributionGroupMember -Identity GroupIdentity [-Credential
Credential ] [-DomainController FullyQualifiedName ]
[-IgnoreDefaultScope {$true | $false}] [-ReadFromDomainController {$true
| $false}] [-ResultSize Size ]
```

Usage

```
Get-DistributionGroupMember -Identity 'CorpSales'
```

You add members to a group using the `Add-DistributionGroupMember` cmdlet. Sample 8-5 provides the syntax and usage.

SAMPLE 8-5 `Add-DistributionGroupMember` cmdlet syntax and usage

Syntax

```
Add-DistributionGroupMember -Identity GroupIdentity [-Member
RecipientIdentity ] [-BypassSecurityGroupManagerCheck {$true | $false}]
[-DomainController FullyQualifiedName ]
```

Usage

```
Add-DistributionGroupMember -Identity 'CorpSales'
-Member 'imaginedlands.com/Sales/April Stewart'
```

You remove members from a group using the `Remove-DistributionGroupMember` cmdlet. Sample 8-6 provides the syntax and usage.

SAMPLE 8-6 `Remove-DistributionGroupMember` cmdlet syntax and usage

Syntax

```
Remove-DistributionGroupMember -Identity GroupIdentity [-Member
RecipientIdentity ] [-BypassSecurityGroupManagerCheck {$true | $false}]
```



[-DomainController FullyQualifiedName]

Usage

Remove-DistributionGroupMember -Identity 'CorpSales'
-Member 'imaginedlands.com/Sales/April Stewart'

Adding and Removing Managers

Group owners are responsible for managing a group. Every group must have at least one owner. To configure a group's managers, follow these steps:

1. In Exchange Admin Center, double-click the group entry. This opens the group's Properties dialog box.
2. On the Ownership page, lists current owners. Click **Add** () to add recipients to the group. In the Select Owners dialog box, select users, groups, or both that should be owners of the group. Select multiple users and groups using the Shift or Ctrl keys.
3. You can remove owners on the Ownership page as well. To remove an owner from a group, select a recipient, then click **Remove** () . When you're finished, click **Save** .



In Exchange Management Shell, you can add or remove group managers using the `-ManagedBy` parameter of the `Set-DistributionGroup` cmdlet. To set this parameter, you must specify the full list of managers for the group by doing the following:

- Add managers by including existing managers and specifying the additional managers when you set the parameter.
- Remove managers by specifying only those who should be managers and excluding those who should not be managers.

If you don't know the current managers of a group, you can list the managers using `Get-DistributionGroup`. You'll need to format the output and examine the value of the `-ManagedBy` property.

Sample 8-7 provides syntax and usage examples for adding and removing group

managers.

SAMPLE 8-7 Adding and removing group managers

Syntax

```
Get-DistributionGroup -Identity GroupIdentity | format-table  
-property ManagedBy
```

```
Set-DistributionGroup -Identity GroupIdentity -ManagedBy GroupManagers
```

Usage

```
Get-DistributionGroup -Identity 'CorpSales' |  
format-table -property ManagedBy
```

```
Set-DistributionGroup -Identity 'CorpSales'  
-ManagedBy 'imaginedlands.com/Sales/Oliver Lee',  
'imaginedlands.com/Users/Jamie Stark'
```

Usage

```
$g = Get-DistributionGroup -Identity 'CorpSales'  
$h = $g.managedby + 'imaginedlands.com/Users/William Stanek'
```

```
Set-DistributionGroup -Identity 'CorpSales'  
-ManagedBy $h
```

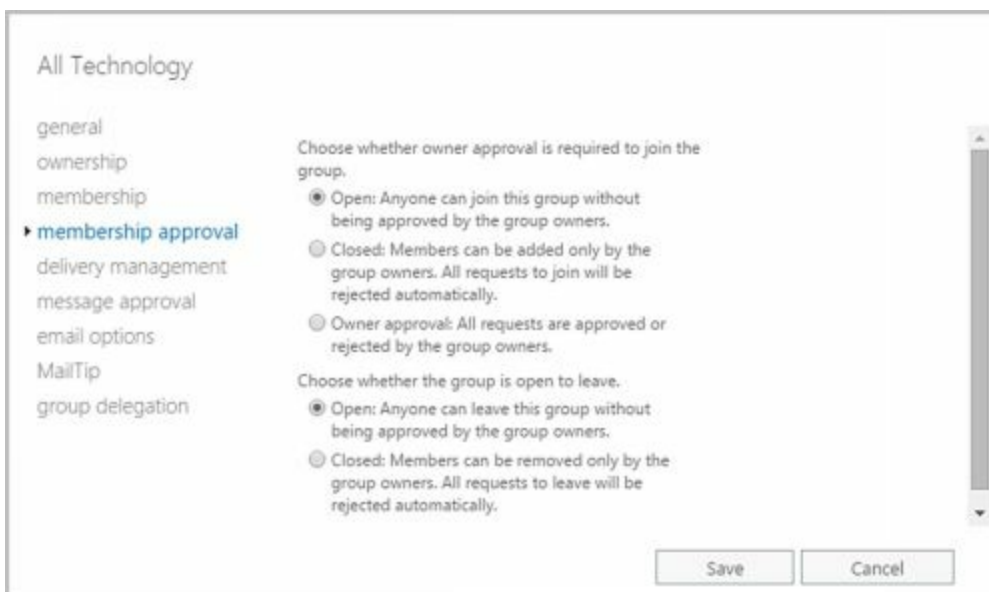
Configuring Member Restrictions and Moderation

Membership in distribution groups can be restricted in several ways. Groups can be open or closed for joining or require group owner approval for joining. Groups can be open or closed for leaving. Groups also can be moderated. With moderated groups, messages are sent to designated moderators for approval before being distributed to members of the group. The only exception is for a message sent by a designated moderator. A message from a moderator is delivered immediately because a moderator has the authority to determine what is and isn't an appropriate message.

To configure member restrictions and moderation, follow these steps:

1. In Exchange Admin Center, double-click the group entry. This opens the group's Properties dialog box.
2. On the Membership Approval page, choose settings for joining the group. The options are:
 - **Open** Anyone can join this group without being approved by the group owners.
 - **Closed** Members can be added only by the group owners. All requests to join will be rejected automatically.
 - **Owner Approval** All requests are approved or rejected by the group owner.

3. Choose settings for leaving the group. The options are:
 - **Open** Anyone can leave this group without being approved by the group owners.
 - **Closed** Members can be removed only by the group owners. All requests to leave will be rejected automatically.
4. The Message Approval page displays the moderation options. To disable moderation, clear the **Messages Sent To This Group Have To Be Approved By A Moderator** check box. To enable moderation, select the **Messages Sent To This Group Have To Be Approved By A Moderator** check box, and then use the options provided to specify group moderators, specify senders who don't require message approval, and configure moderation notifications.
5. Click **Save** to apply your changes.



In Exchange Management Shell, you manage distribution group settings using Set-DistributionGroup. You configure member restrictions for joining a group using the -MemberJoinRestriction parameter and configure member restrictions for leaving a group using the -MemberDepartRestriction parameter. If you want to check the current restrictions, you can do this using Get-DistributionGroup. You'll need to format the output and examine the values of the -MemberJoinRestriction property, the -MemberDepartRestriction property, or both.

Sample 8-8 provides syntax and usage examples for configuring member restrictions.

SAMPLE 8-8 Configuring member restrictions for groups

Syntax

```
Get-DistributionGroup -Identity GroupIdentity | format-table -property
Name, MemberJoinRestriction, MemberDepartRestriction
```

```
Set-DistributionGroup -Identity GroupIdentity
[-MemberJoinRestriction <Closed | Open | ApprovalRequired>]
[-MemberDepartRestriction <Closed | Open | ApprovalRequired>]
```


Usage

```
Get-DistributionGroup -Identity 'AllMarketing' |  
format-table -property Name, MemberJoinRestriction,  
MemberDepartRestriction
```

```
Set-DistributionGroup -Identity 'AllMarketing'  
-MemberJoinRestriction 'Closed' -MemberDepartRestriction 'Closed'
```

Set-DistributionGroup parameters for configuring moderation include -ModerationEnabled, -ModeratedBy, -BypassModerationFromSendersOrMembers, and -SendModerationNotifications. You enable or disable moderation by using -ModerationEnabled. If moderation is enabled, you can do the following:

- Designate moderators using -ModeratedBy.
- Specify senders who don't require message approval by using -BypassModerationFromSendersOrMembers.
- Configure moderation notifications using -SendModerationNotifications.

Sample 8-9 provides syntax and usage examples for configuring moderation.

SAMPLE 8-9 Configuring moderation for groups

Syntax

```
Get-DistributionGroup -Identity GroupIdentity | format-table -property  
Name, ModeratedBy, BypassModerationFromSendersOrMembers,  
SendModerationNotifications
```

```
Set-DistributionGroup -Identity GroupIdentity  
[-ModeratedBy Moderators ] [-ModerationEnabled <$true | $false>]  
[-BypassModerationFromSendersOrMembers Recipients ]  
[-SendModerationNotifications <Never | Internal | Always>]
```

Usage

```
Get-DistributionGroup -Identity 'AllMarketing' |  
format-table -property Name, ModeratedBy,  
BypassModerationFromSendersOrMembers, SendModerationNotifications
```

```
Set-DistributionGroup -Identity 'AllMarketing'  
-ModerationEnabled $true -Moderators 'AprilC'  
-SendModerationNotifications 'Internal'
```

Working with Dynamic Distribution Groups

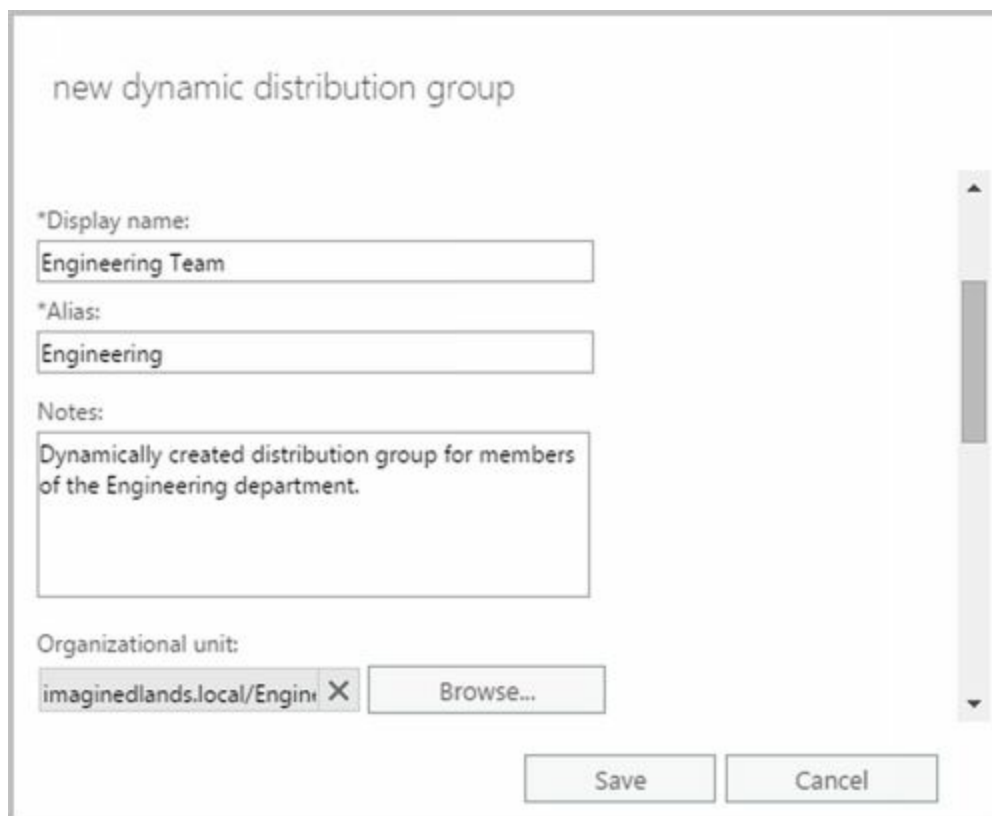
Just as there are tasks that apply only to security and standard distribution groups, there are also tasks that apply only to dynamic distribution groups. These tasks are discussed in this section.

Creating Dynamic Distribution Groups

With dynamic distribution groups, group membership is determined by the results of an LDAP query. You can create a dynamic distribution group and define the query parameters by completing the following steps:

1. In Exchange Admin Center, select **Recipients** in the Navigation menu and then select **Groups**.

2. Click **New** () and then select **Dynamic Distribution Group**. This opens the New Dynamic Distribution Group dialog box, shown in Figure 8-4.



The screenshot shows a dialog box titled "new dynamic distribution group". It contains the following fields and controls:


- *Display name:** A text box containing "Engineering Team".
- *Alias:** A text box containing "Engineering".
- Notes:** A text area containing "Dynamically created distribution group for members of the Engineering department."
- Organizational unit:** A text box containing "imaginedlands.local/Engin" with a close button (X) and a "Browse..." button.
- Buttons:** "Save" and "Cancel" buttons at the bottom.

FIGURE 8-4 Configuring the basic settings for the dynamic distribution group.

3. In the Display Name text box, type a display name for the group. Group names aren't case-sensitive and can be up to 64 characters long. Keep in mind that group naming policy doesn't apply to administrators creating distribution groups in Exchange Admin Center.
4. Like users, groups have Exchange aliases. Enter an alias. The Exchange alias is used to set the group's SMTP e-mail address. Exchange Server uses the SMTP address for receiving messages.

5. With on-premises Exchange, the group account is created in the default user container, which typically is the Users container. To create the group in a specific organizational unit instead, click **Browse** to the right of the Organizational Unit text box. In the Select Organizational Unit dialog box, choose the location where you want to store the account and then click **OK**.

NOTE With Exchange 2016, the organizational unit you specify is simply the storage container. Thus, unlike Exchange 2010, the selection is not used to scope or filter the LDAP query.

6. Group owners are responsible for managing groups. Unlike standard distribution groups, dynamic distribution groups don't need to be assigned an owner. If you want to specify an owner, under Owner, click **Add** (). In the Select Owner dialog box, select the user or group that should have management responsibility for the group.
7. Specify the recipients to include in the group. To allow any recipient type to be a member of the group, select **All Recipient Types**. Otherwise, choose **Only The Following Recipient Types** and then choose the types of recipients to include in the dynamic distribution group.
8. Membership in the group is determined by the rules you define. To define a rule, click **Add A Rule** and set the filter conditions. The following types of conditions as well as conditions for custom attributes are available:

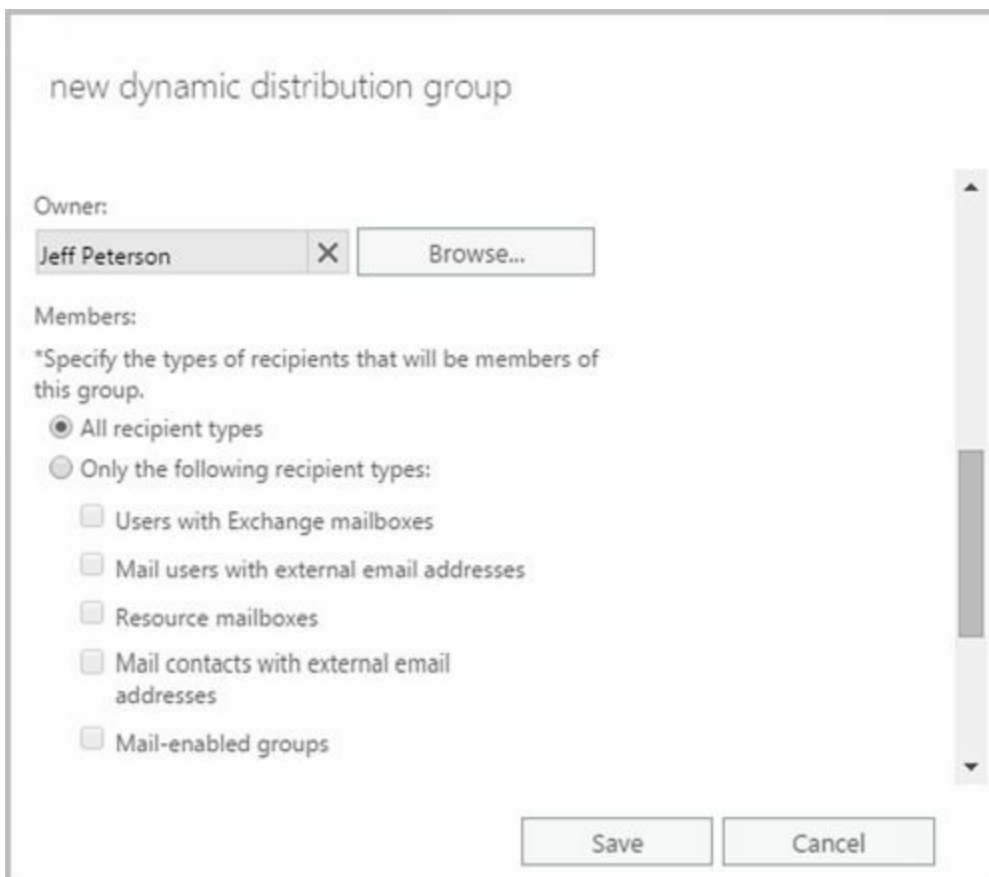


FIGURE 8-5 Configuring ownership and membership settings for the dynamic distribution group.

- **Recipient Container** Filters recipients based on where the related account is stored in Active Directory. Selecting this option displays the Select An Organizational Unit dialog box. Click the container where the recipients are stored, such as Users or an organizational unit, and then click OK.
- **State Or Province** Filters recipients based on the value of the State/Province text box on the Contact Information page in the related Properties dialog box. Selecting this option displays the Specify Words Or Phrases dialog box. Type a state or province identifier to use as a filter condition and then press Enter or click Add. Repeat as necessary, and then click OK.
- **Department** Filters recipients based on the value of the Department text box on the Organization page in the related Properties dialog box. Selecting this option displays the Specify Words Or Phrases dialog box. Type a department name to use as a filter condition and then press Enter or click Add. Repeat as necessary, and then click OK.
- **Company** Filters recipients based on the value of the Company text box on the Organization page in the related Properties dialog box. Selecting this option displays the Specify Words Or Phrases dialog box. Type a company name to use as a filter condition and then press Enter or click Add. Repeat as necessary, and then click OK.

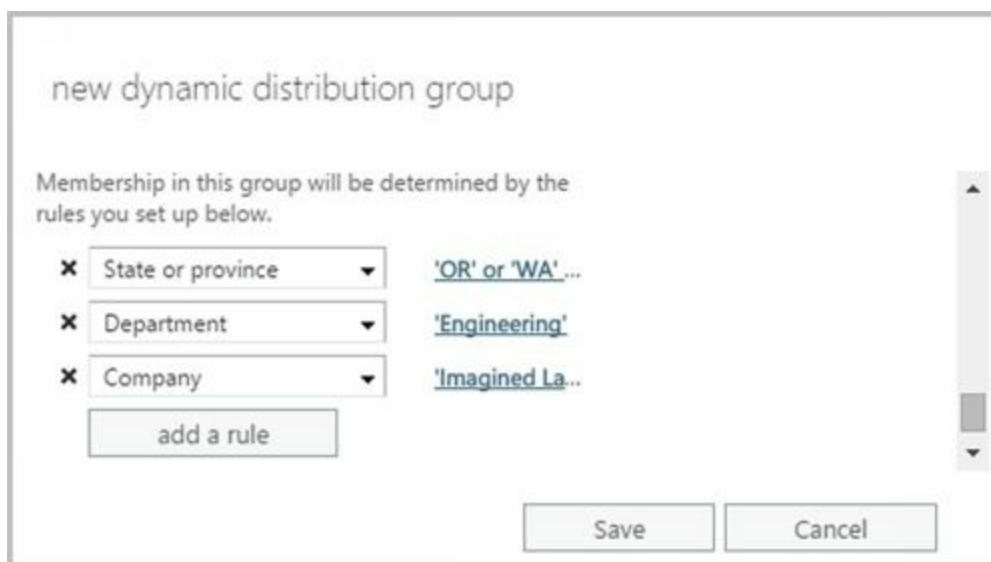


FIGURE 8-6 Setting the filter conditions.

IMPORTANT Although each rule acts as an OR condition for matches on specified values, the rules are aggregated as AND conditions. This means that a user that matches one of the values in a rule passes that filter but must be a match for all the rules to be included in the group. For example, if you were to define a state rule for Oregon, California, or Washington and a department rule for Technology, only users who are in Oregon, California, or Washington *and* in the Technology department match the filter and are included as members of the group.

9. Click **Save** to create the group. If an error occurs during group creation, the related group will not be created. You need to correct the problem before you can complete this procedure.
10. Creating the group isn't the final step. Afterward, you might want to do the following:

- Set message size restrictions for messages mailed to the group.
- Limit users who can send to the group.
- Change or remove default email addresses.
- Add more email addresses.

In Exchange Management Shell, you can create a dynamic distribution group using the `New-DynamicDistributionGroup` cmdlet. Sample 8-10 provides the syntax and usage.

SAMPLE 8-10 `New-DynamicDistributionGroup` cmdlet syntax and usage

Syntax

```
New-DynamicDistributionGroup -Name ExchangeName
-IncludedRecipients <None, MailboxUsers, MailContacts, MailGroups,
Resources, AllRecipients> [-Alias ExchangeAlias ]
[-DisplayName DisplayName ] [-OrganizationalUnit OUName ]
[-ConditionalCompany CompanyNameFilter1 , CompanyNameFilter2 ,...]
[-ConditionalCustomAttribute X Value1 , Value2,... ]
[-ConditionalDepartment DeptNameFilter1 , DeptNameFilter2 , ... ]
[-ConditionalStateOrProvince StateNameFilter1 , StateNameFilter2 , ...]
[-RecipientContainer ApplyFilterContainer] {AddtlParams}
```

```
New-DynamicDistributionGroup -Name ExchangeName -RecipientFilter Filter
[-Alias ExchangeAlias ] [-DisplayName DisplayName ] [-OrganizationalUnit
OUName ] [-RecipientContainer ApplyFilterContainer ] {AddtlParams}
```

{AddtlParams}

```
[-ArbitrationMailbox ModeratorMailbox ] [-DomainController
FullyQualifiedName ] [-ExternalDirectoryObjectId ObjectId ]
[-ModeratedBy Moderators ] [-ModerationEnabled <$true | $false>]
[-Organization OrgName ] [-PrimarySmtpAddress SmtpAddress ]
[-SendModerationNotifications <Never | Internal | Always>]
```

Usage

```
New-DynamicDistributionGroup -Name 'CrossSales'
-OrganizationalUnit 'imaginedlands.com/Users' -DisplayName
'CrossSales' -Alias 'CrossSales'
-IncludedRecipients 'MailboxUsers, MailContacts, MailGroups'
-ConditionalCompany 'Imagined Lands'
-ConditionalDepartment 'Sales','Marketing'
-ConditionalStateOrProvince 'Washington','Oregon','California'
-RecipientContainer 'imaginedlands.com'
```

Changing Query Filters and Filter Conditions

With dynamic distribution groups, the filter conditions determine the exact criteria that must be met for a recipient to be included in the dynamic distribution group. You can modify the filter conditions by completing the following steps:

1. In Exchange Admin Center, double-click the dynamic distribution group entry. This opens the group's Properties dialog box.
2. On the Membership page, use the Specify The Types Of Recipients options to specify the types of recipients to include in the query. Select either **All Recipient Types** or select **Only The Following Recipient Types**, and then select the types of recipients.
3. The Membership page lists the current conditions. The following types of conditions as well as conditions for custom attributes are available:
 - **State Or Province** Filters recipients based on the value of the State/Province text box on the Contact Information page in the related Properties dialog box. Click the related Enter Words link. In the Specify Words Or Phrases dialog box, type a state or province identifier to use as a filter condition and then press Enter or click Add. Repeat as necessary, and then click OK.
 - **Department** Filters recipients based on the value of the Department text box on the Organization page in the related Properties dialog box. Click the related Enter Words link. In the Specify Words Or Phrases dialog box, type a department name to use as a filter condition and then press Enter or click Add. Repeat as necessary, and then click OK.
 - **Company** Filters recipients based on the value of the Company text box on the Organization page in the related Properties dialog box. Click the related Enter Words link. In the Specify Words Or Phrases dialog box, type a company name to use as a filter condition and then press Enter or click Add. Repeat as necessary, and then click OK.
4. Click **Save** to apply the changes.

The screenshot shows the 'Engineering Team' Properties dialog box, specifically the 'membership' tab. On the left, a navigation pane lists various settings: general, ownership, membership (selected), delivery management, message approval, email options, MailTip, and group delegation. The main area is titled 'Members:' and contains the following options:

- *Specify the types of recipients that will be members of this group.
 - All recipient types
 - Only the following recipient types:
 - Users with Exchange mailboxes
 - Mail users with external email addresses
 - Resource mailboxes
 - Mail contacts with external email addresses
 - Mail-enabled groups

Below these options, a note states: 'Membership in this group will be determined by the rules you set up below.' There are four rules listed, each with a dropdown menu and a value:

- Recipient container: Engineering
- State or province: 'CA' or 'WA' or 'OR'
- Company: 'Imagined Lands'
- Department: 'Engineering'

At the bottom of the rules list is an 'add a rule' button. At the very bottom of the dialog box are 'Save' and 'Cancel' buttons.

Designating an Expansion Server

When there are potentially hundreds or thousands of members, dynamic distribution

groups are inefficient and can require a great deal of processing to complete. This is why Exchange 2016 shifts the processing requirements from Global Catalog servers to dedicated expansion servers. However, the routing destination is the ultimate destination for a message. A distribution group expansion server is the routing destination when a distribution group has a designated expansion server that's responsible for expanding the membership list of the group. A distribution group expansion server is always an Exchange 2016 Mailbox server, an Exchange 2013 Mailbox server or an Exchange 2010 Hub Transport server.

Each routing destination has a delivery group, which is a collection of one or more transport servers that are responsible for delivering messages to that routing destination. When the routing destination is a distribution group expansion server, the delivery group may contain Exchange 2016 Mailbox servers, Exchange 2013 Mailbox servers, and Exchange 2010 Hub Transport servers.

How the message is routed depends on the relationship between the source transport server and the destination delivery group. If the source transport server is in the destination delivery group, the routing destination itself is the next hop for the message. The message is delivered by the source transport server to the mailbox database or connector on a transport server in the delivery group.

On the other hand, if the source transport server is outside the destination delivery group, the message is relayed along the least-cost routing path to the destination delivery group. In a complex Exchange organization, a message may be relayed to other transport servers along the least-cost routing path or relayed directly to a transport server in the destination delivery group.

REAL WORLD Keep in mind that when a distribution group expansion server is the routing destination, the distribution group is already expanded when a message reaches the routing stage of categorization on the distribution group expansion server. Therefore, the routing destination from the distribution group expansion server is always a mailbox database or a connector.

By default, Exchange 2016 uses the closest Exchange server that has the Mailbox server role installed as the dedicated expansion server. Because routing destinations and delivery groups can also include Exchange 2013 Mailbox servers and Exchange 2010 Hub Transport servers in mixed environments, Exchange 2013 and Exchange 2010 servers could perform distribution group expansion in mixed Exchange organizations.

In some cases, you might want to explicitly specify the dedicated expansion server to handle expansion processing for some or all of your dynamic distribution groups. A key reason for this is to manage where the related processing occurs and in this way shift the processing overhead from other servers to this specified server. You can specify a dedicated expansion server using the `-ExpansionServer` parameter of the `Set-DynamicDistributionGroup` cmdlet.

Modifying Dynamic Distribution Groups Using Cmdlets

In Exchange Management Shell, you can use the `Get-DynamicDistributionGroup` cmdlet to get information about dynamic distribution groups and modify their associated filters and conditions using the `Set-DynamicDistributionGroup` cmdlet.

Sample 8-11 provides the syntax and usage for the `Get-DynamicDistributionGroup` cmdlet.

SAMPLE 8-11 `Get-DynamicDistributionGroup` cmdlet syntax and usage

Syntax

```
Get-DynamicDistributionGroup [-Identity GroupIdentify | -Anr Name
| -ManagedBy Managers ]
[-AccountPartition PartitionID ] [-Credential Credential ]
[-DomainController FullyQualifiedName ] [-Filter FilterString ]
[-IgnoreDefaultScope {$true | $false}] [-Organization OrgName ]
[-OrganizationalUnit OUName ] [-ReadFromDomainController {$true | $false}]
[-ResultSize Size ] [-SortBy Value ]
```

Usage

```
Get-DynamicDistributionGroup -Identity 'CrossSales'
```

Sample 8-12 provides the syntax and usage for the `Set-DynamicDistributionGroup` cmdlet.

SAMPLE 8-12 `Set-DynamicDistributionGroup` cmdlet syntax and usage

Syntax

```
Set-DynamicDistributionGroup -Identity GroupIdentity
[-Alias NewAlias ] [-AcceptMessagesOnlyFrom Recipients ]
[-AcceptMessagesOnlyFromDLMembers Recipients ]
[-AcceptMessagesOnlyFromSendersOrMembers Recipients ]
[-ArbitrationMailbox ModeratorMailbox ]
[-BypassModerationFromSendersOrMembers Recipients ]
[-ConditionalCompany Values ] [-ConditionalDepartment Values ]
[-ConditionalCustomAttribute X Values]
[-ConditionalStateOrProvince Values ] [-CreateDTMFMap <$true | $false>]
[-DisplayName Name ] [-DomainController DCName]
[-EmailAddresses ProxyAddress]
[-EmailAddressPolicyEnabled <$false|$true>]
[-ExpansionServer Server ] [-ForceUpgrade <$false|$true>]
[-ExtensionCustomAttribute X Value1 , Value2,... ]
[-GrantSendOnBehalfTo Mailbox]
[-HiddenFromAddressListsEnabled <$false|$true>]
[-IgnoreDefaultScope {$true | $false}]
[-IncludedRecipients <None, MailboxUsers, MailContacts, MailGroups,
Resources, AllRecipients>] [-MailTip String ]
[-MailTipTranslations Locale : TipString, Locale : TipString, ... ]
[-ManagedBy Managers ] [-MaxReceiveSize Size ] [-MaxSendSize Size]
[-ModeratedBy Moderators ] [-ModerationEnabled <$true | $false>]
```


[-Name **Name**] [-Notes **Value**] [-PhoneticDisplayName **PhName**]
[-PrimarySmtpAddress **SmtpAddress**]
[-RecipientContainer **OUName**] [-RecipientFilter **String**]
[-RejectMessagesFrom **Recipients**]
[-RejectMessagesFromDLMembers **Recipients**]
[-RejectMessagesFromSendersOrMembers **Recipients**]
[-ReportToManagerEnabled <\$false|\$true>]
[-ReportToOriginatorEnabled <\$false|\$true>]
[-RequireSenderAuthenticationEnabled <\$false|\$true>]
[-SendModerationNotifications <Never | Internal | Always>]
[-SendOofMessageToOriginatorEnabled <\$false|\$true>]
[-SimpleDisplayName **Name**] [-UMDtmfMap **Values**]
[-WindowsEmailAddress **SmtpAddress**]

Usage

```
Set-DynamicDistributionGroup -Identity 'CrossSales'  
-IncludedRecipients 'AllRecipients'  
-ConditionalCompany 'Imagined Lands'  
-ConditionalDepartment 'Sales','Accounting'  
-ConditionalStateOrProvince 'Washington','Idaho','Oregon'  
-RecipientContainer 'imaginedlands.com'
```

Usage

```
Set-DynamicDistributionGroup -Identity 'CrossSales'  
-ForceUpgrade $true
```

Usage

```
Set-DynamicDistributionGroup -Identity 'CrossSales'  
-ExpansionServer 'CorpSvr127'
```

Previewing Dynamic Distribution Group Membership

You can preview a dynamic distribution group to confirm its membership and determine how long it takes to return the query results. The specific actions you take depend on the following factors:

- In some cases, membership isn't what you expected. If this happens, you need to change the query filters, as discussed earlier.
- In other cases, it takes too long to execute the query and return the results. If this happens, you might want to rethink the query parameters and create several query groups.

You can quickly determine how many recipients are in the group by checking how many recipients received the last message sent to the group. One way to do this is to follow these steps:

1. In Exchange Admin Center, select the dynamic distribution group entry.

2. In the details pane, look under Membership to see the number of recipients who received the last message sent to the group.

In Exchange Management Shell, you can determine the exact membership of a dynamic distribution group by getting the dynamic group and then using the associated recipient filter to list the members. Consider the following example:

```
$Members = Get-DynamicDistributionGroup "TechTeam"  
Get-Recipient -RecipientPreviewFilter $Members.RecipientFilter
```

In this example, `Get-DynamicDistributionGroup` stores the object for the TechTeam group in the `$Members` variable. Then `Get-Recipient` lists the recipients that match the recipient filter on this object. Note that the Exchange identifier can be the display name or alias for the group.

Other Essential Tasks for Managing Groups

Previous sections covered tasks that were specific to a type of group. As an Exchange administrator, you'll need to perform many additional group management tasks. These essential tasks are discussed in this section.

Changing a Group's Name Information

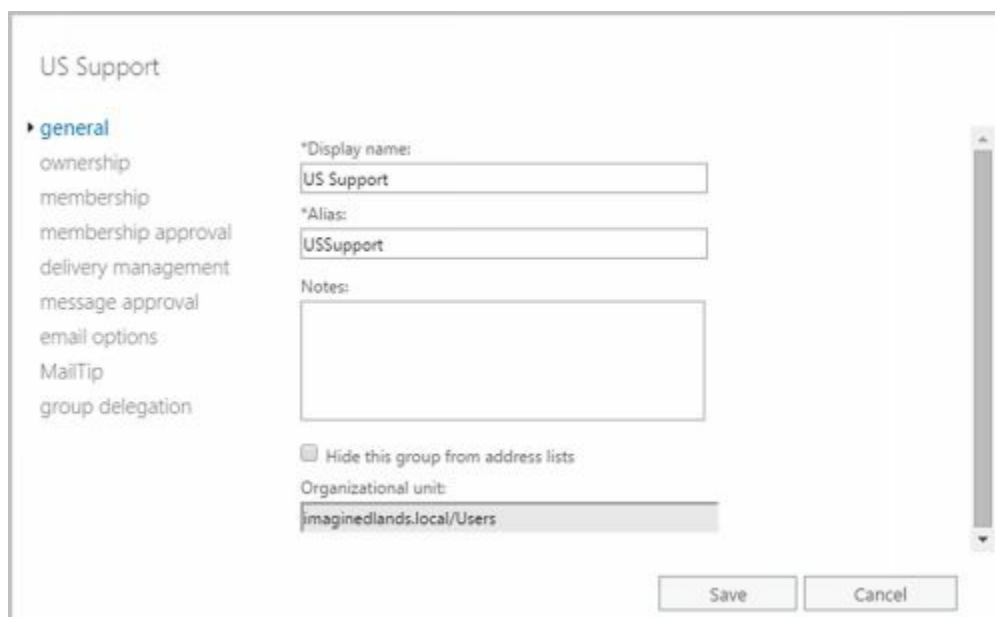
Each mail-enabled group has a display name, an Exchange alias, and one or more email addresses associated with it. The display name is the name that appears in address lists. The Exchange alias is used to set the email addresses associated with the group.

Whenever you change a group's naming information, new email addresses can be generated and set as the default addresses for SMTP. These email addresses are used as alternatives to email addresses previously assigned to the group. To learn how to change or delete these additional email addresses, see the "Changing, Adding, or Deleting a Group's Email Addresses" section later in this chapter.

To change the group's Exchange name details, complete the following steps:

1. In Exchange Admin Center, double-click the group entry. This opens the group's Properties dialog box.
2. On the General page, the first text box shows the display name of the group. If necessary, type a new display name.
3. The Alias text box shows the Exchange alias. If necessary, type a new alias. Click **Save** .

NOTE When you change a group's display name, you give the group a new label. Changing the display name doesn't affect the SID, which is used to identify, track, and handle permissions independently from group names.



The screenshot shows the 'US Support' group's Properties dialog box in the Exchange Admin Center. The 'general' tab is selected. The 'Display name' field contains 'US Support'. The 'Alias' field contains 'USSupport'. The 'Notes' field is empty. There is a checkbox for 'Hide this group from address lists' which is unchecked. The 'Organizational unit' field contains 'imaginedlands.local/Users'. At the bottom, there are 'Save' and 'Cancel' buttons.


Changing, Adding, or Deleting a Group's Email Addresses

When you create a mail-enabled group, default email addresses are created for SMTP. Any time you update the group's Exchange alias, new default email addresses can be created. The old addresses aren't deleted, however; they remain as alternative email addresses for the group.

To change, add, or delete a group's email addresses, follow these steps:

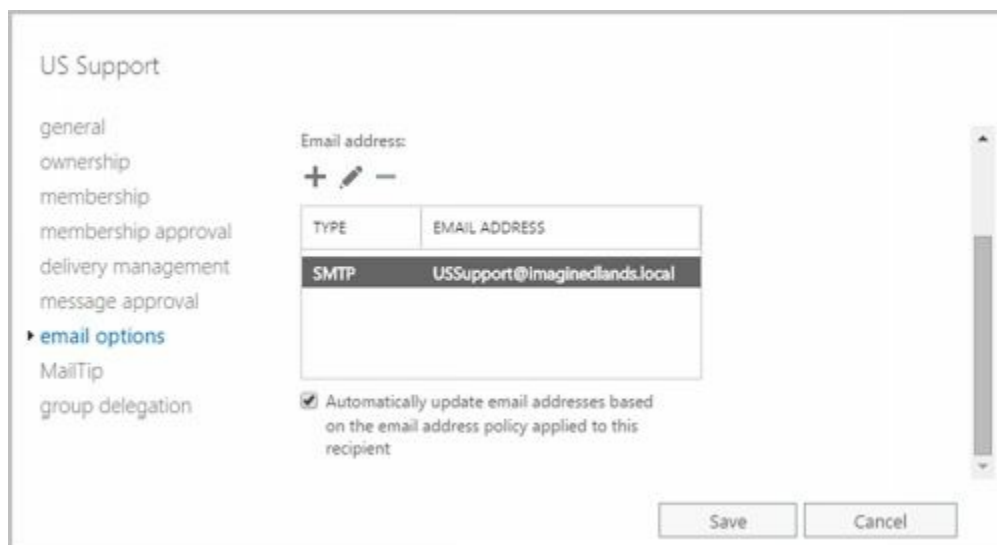
1. In Exchange Admin Center, double-click the group entry. This opens the group's Properties dialog box.
2. On the Email Options page, use the following techniques to manage the group's email addresses:

- **Create a new SMTP address** Click Add (). In the New Email Address dialog box, SMTP is selected as the address type by default. Enter the email address, and then click OK.

- **Create a custom address** Click Add (). In the New Email Address dialog box, select Custom Address Type. Enter a prefix that identifies the type of email address, and then enter the associated address. Click OK.

TIP Use SMTP as the address type for standard Internet email addresses. For custom address types, such as X.400, you must manually enter the address in the proper format.

- **Set a new Reply To Address** Double-click the address that you want to use as the primary SMTP address. Select Make This The Reply Address, and then click OK. (Exchange Online Only)
- **Edit an existing address** Double-click the address entry. Modify the settings in the Address dialog box, and then click OK.
- **Delete an existing address** Select the address, and then click Remove.



Sample 8-13 provides syntax and usage examples for configuring a group's primary SMTP email address. If email address policy is enabled, you won't be able to update the email address unless you set -EmailAddressPolicyEnabled to \$false.

SAMPLE 8-13 Configuring a group's primary SMTP email address

Syntax

```
Get-DistributionGroup-Identity GroupIdentity | format-list -property  
Name, EmailAddresses, PrimarySmtpAddress
```

```
Set-DistributionGroup-Identity GroupIdentity  
-PrimarySmtpAddress SmtpAddress -EmailAddressPolicyEnabled $false
```

Usage

```
Get-DistributionGroup -Identity 'AllSales' | format-list -property  
Name, EmailAddresses, PrimarySmtpAddress
```

```
Set-DistributionGroup -Identity 'AllSales'  
-PrimarySmtpAddress allsales@imaginedlands.com  
-EmailAddressPolicyEnabled $false
```

Hiding Groups from Exchange Address Lists

By default, any mail-enabled security group or other distribution group that you create is shown in Exchange address lists, such as the global address list. If you want to hide a group from the address lists, follow these steps:

1. In Exchange Admin Center, double-click the group entry. This opens the group's Properties dialog box.
2. On the General page, select the **Hide This Group From Address Lists** check box. Click **OK**.



NOTE When you hide a group, it isn't listed in Exchange address lists. However, if a user knows the name of a group, he or she can still use it in the mail client. To prevent users from sending to a group, you must set message restrictions, as discussed in the next section, "Setting Usage Restrictions on Groups."

TIP Hiding group membership is different from hiding the group itself. In Outlook, users can view the membership of groups. In Exchange Server 2016, you

cannot prevent viewing the group membership. In addition, membership of dynamic distribution groups is not displayed in global address lists because it is generated only when mail is sent to the group.

In Exchange Management Shell, you can return a list of groups hidden from address lists using either of the following commands:

```
Get-DistributionGroup -filter {HiddenFromAddressListsEnabled -eq $true}
```

```
Get-DistributionGroup | where {$_.HiddenFromAddressListsEnabled -eq $true}
```

Setting Usage Restrictions on Groups

Groups are great resources for users in an organization. They let users send mail quickly and easily to other users in their department, business unit, or office. However, if you aren't careful, people outside the organization could use groups as well. Would your boss like it if spammers sent unsolicited email messages to company employees through your distribution lists? Probably not—and you'd probably be sitting in the hot seat, which would be uncomfortable, to say the least.

To prevent unauthorized use of mail-enabled groups, groups are configured by default to accept mail only from authenticated users so that only senders inside an organization can send messages to groups. An authenticated user is any user accessing the system through a logon process. It does not include anonymous users or guests. If you use the default configuration, any message from a sender outside the organization is rejected. Off-site users will need to log on to Exchange before they can send mail to groups, which might present a problem for users who are at home or travelling.

REAL WORLD If you have users who telecommute or send email from home using a personal account, you might be wondering how these users can send mail with a restriction that allows only senders inside the organization to send messages to the group. What I've done in the past is create a group called `OffsiteEmailUsers` and then added this as a group that can send mail to my mail-enabled groups. The `OffsiteEmailUsers` group contains separate mail-enabled contacts for each authorized off-site email address. Alternatively, users could simply log on using MAPI over HTTP, Outlook Anywhere (RPC over HTTP), Outlook Web App, or Exchange ActiveSync and send mail to the group; this is an approach that doesn't require any special groups with permissions to be created or maintained.

Alternatively, you can allow senders inside and outside the organization to send email to a group. This setting allows unrestricted access to the group, so anyone can send messages to the group. However, this exposes the group to spam from external mail accounts.

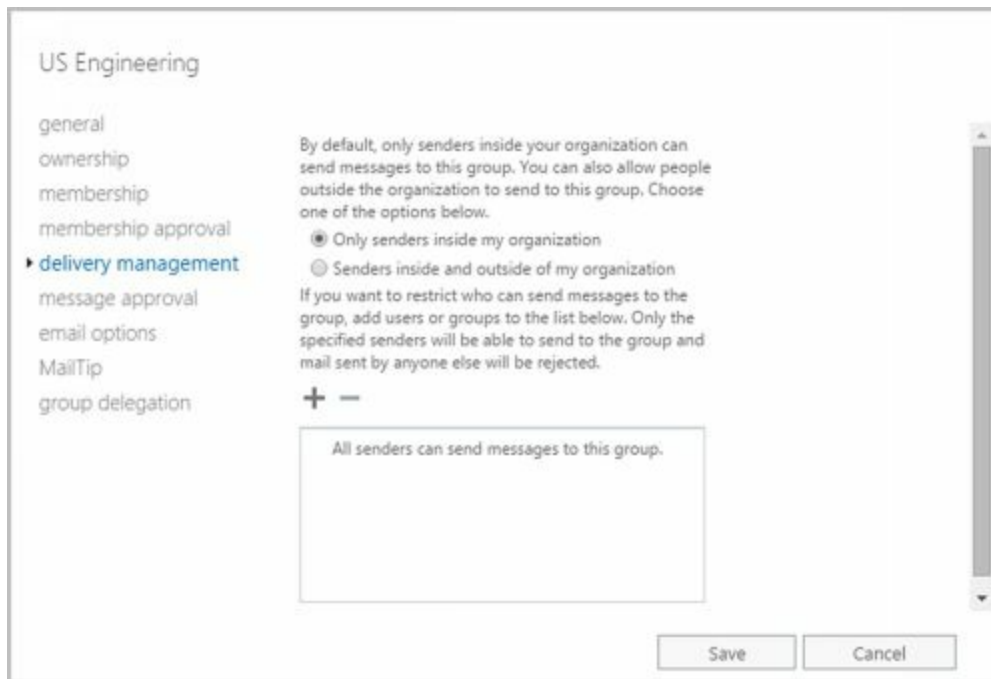
Another way to prevent unauthorized use of mail-enabled groups is to specify that only certain users or members of a particular group can send messages to the group. For example, if you create a group called `AllEmployees`, of which all company employees are members, you can specify that only the members of `AllEmployees` can send


messages to the group. You do this by specifying that only messages from members of AllEmployees are acceptable.

To prevent mass spamming of other groups, you can set the same restriction. For example, if you have a group called Technology, you could specify that only members of AllEmployees can send messages to that group.

You can set or remove usage restrictions by completing the following steps:

1. In Exchange Admin Center, double-click the group entry. In the Properties dialog box for the group, select the **Delivery Management** page.



2. To ensure that messages are accepted only from authenticated users, select **Only Senders Inside My Organization**.
3. To accept messages from all email addresses, select **Senders Inside And Outside Of My Organization**.
4. To restrict senders, specify that messages only from the listed users, contacts, or groups be accepted. To do this, click **Add** () to display the Select Allowed Senders dialog box. Select a recipient, and then click **OK**. Repeat as necessary.

TIP You can select multiple recipients at the same time. To select multiple recipients individually, hold down the Ctrl key and then click each recipient that you want to select. To select a continuous sequence of recipients, select the first recipient, hold down the Shift key, and then click the last recipient.

5. Click **Save**.

Creating Moderated Groups

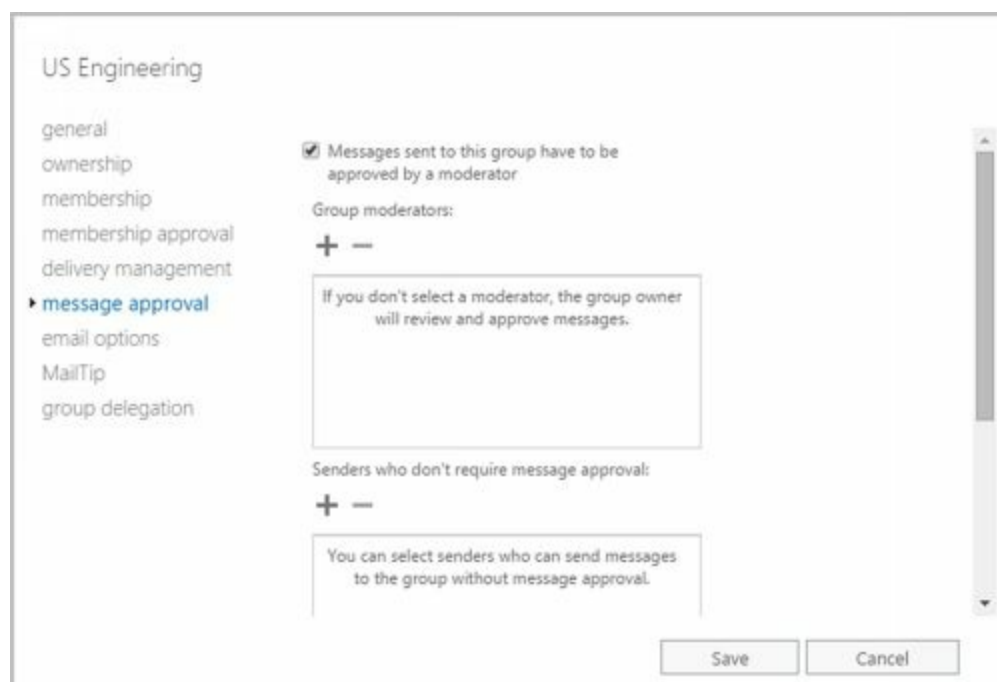
By default, senders don't require approval for their messages to be sent to all members

of a group. Sometimes though you'll want to appoint moderators who must approve messages before they are sent to all members of the group. If you enable moderation but don't specify a moderator or moderators, the group owner is responsible for reviewing and approving messages. When moderation is enabled, you also can specify users who don't require approval for their messages to be sent to all members of the group.

To see how moderation could be used, consider the following example. A project team is set up to work on a restricted project. The team leader wants a moderated group for the project team so that she must review and approve all messages sent to the group before they are sent to members of the team. As the moderator, the team leader's messages don't require approval and are sent directly to all members of the group.

To configure moderation for a group, complete the following steps:

1. In Exchange Admin Center, double-click the group name to open the Properties dialog box for the group.
2. On the Message Approval page, do one of the following:
 - To enable moderation, select **Messages Sent To This Group Have To Be Approved By A Moderator**. Next, use the options provided to specify moderators and senders who don't required message approval.
 - To disable moderation, clear **Messages Sent To This Group Have To Be Approved By A Moderator**. Click **Save** and then skip the rest of the steps.



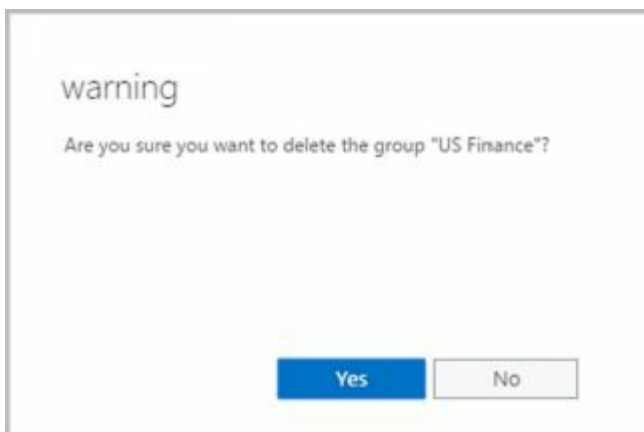
3. Use the Group Moderators options to add moderators. If there are any senders who don't require message approval, add these as well using the options provided.
4. If a message addressed to the group isn't approved, the message isn't distributed to members of the group, and all users receive a nondelivery report (NDR) by default whether they are inside or outside the organization. Alternatively, you can notify only senders in your organization when their messages aren't approved or you can disable notification completely.

5. Click Save .



Deleting Groups

If you are an owner of a group, you can delete it. Deleting a group removes it permanently. After you delete a security group, you can't create a security group with the same name and automatically restore the permissions that the original group was assigned because the SID for the new group won't match the SID for the old group. You can reuse group names, but remember that you'll have to re-create all permissions settings.



You cannot delete built-in groups in Windows. In Exchange Admin Center, you can remove other types of groups by selecting them and clicking Delete. When prompted, click Yes to delete the group. If you click No, Exchange Admin Center will not delete the group.

In Exchange Management Shell, only a group's manager or other authorized user can remove a group. Use the Remove-DistributionGroup cmdlet to remove distribution groups, as shown in Sample 8-14.

SAMPLE 8-14 Remove-DistributionGroup cmdlet syntax and usage

Syntax

```
Remove-DistributionGroup -Identity GroupIdentity  
[-BypassSecurityGroupManagerCheck {$true | $false}]  
[-DomainController FullyQualifiedName ]  
[-IgnoreDefaultScope {$true | $false}]
```

Usage

```
Remove-DistributionGroup -Identity 'imaginedlands.com/Users/AllSales'
```

To remove dynamic distribution groups, you can use the Remove-DynamicDistributionGroup cmdlet. Sample 8-15 shows the syntax and usage.

SAMPLE 8-15 Remove-DynamicDistributionGroup cmdlet syntax and usage

Syntax

```
Remove-DynamicDistributionGroup -Identity GroupIdentity  
[-DomainController FullyQualifiedName ]  
[-IgnoreDefaultScope {$true | $false}]
```

Usage

```
Remove-DynamicDistributionGroup -Identity 'CrossSales'
```

Chapter 9. Managing Addresses Online and Offline

Email addresses and other contact information stored in the Exchange organization are available via online address lists and offline address books. Address lists are collections of recipients in an Exchange organization that are selectable in the address book of client applications. Exchange can be configured to distribute copies of address books to authorized users so that address lists are accessible when working offline.

Managing Online Address Lists

You use address lists to organize recipients by department, business unit, location, type, and other criteria. The default address lists that Exchange Server creates, as well as any new address lists that you create, are available to the user community based on their view of the global address list. Users can navigate these address lists to find recipients to whom they want to send messages.

Using Default Address Lists

During setup, Exchange Server creates a number of default address lists that are selectable in the address book of client applications, including the following:

- **Default Global Address List** Lists all mail-enabled users, contacts, and groups in the organization.
- **Default Offline Address Book** Provides an address list for viewing offline that contains information on all mail-enabled users, contacts, and groups in the organization.
- **All Contacts** Lists all mail-enabled contacts in the organization.
- **All Groups** Lists all mail-enabled groups in the organization.
- **All Rooms** Lists all resource mailboxes for rooms.
- **Public Folders** Lists all public folders in the organization.
- **All Users** Lists all mail-enabled users in the organization.

IMPORTANT Generally, whenever you specify address list paths in Exchange Management Shell, you must reference their position relative to the root container. The root container is identified as \. If the address list name contains spaces, you also must enclose the address list path in quotes. Thus, you reference the Default Address List as '\Default Address List' and All Rooms as '\All Rooms'.

The most commonly used address lists are the global address list and the offline address book. In Exchange Admin Center for your on-premises organization, you access online address lists and offline address books by selecting Organization in the Navigation menu and then selecting Address Lists. As Figure 9-1 shows, the main pane shows each address list by name and up-to-date status. If an address list isn't up-to-date, you can click the Update option to update it.

IMPORTANT Any address list created using the shell should be managed only with the shell. Address lists created with the GUI can be managed with either the GUI or the shell. That said, Microsoft recommends that you manage address lists from the shell whenever the list contains several thousand or more recipients. The reason for this is that Exchange Admin Center will be locked until the task is completed.

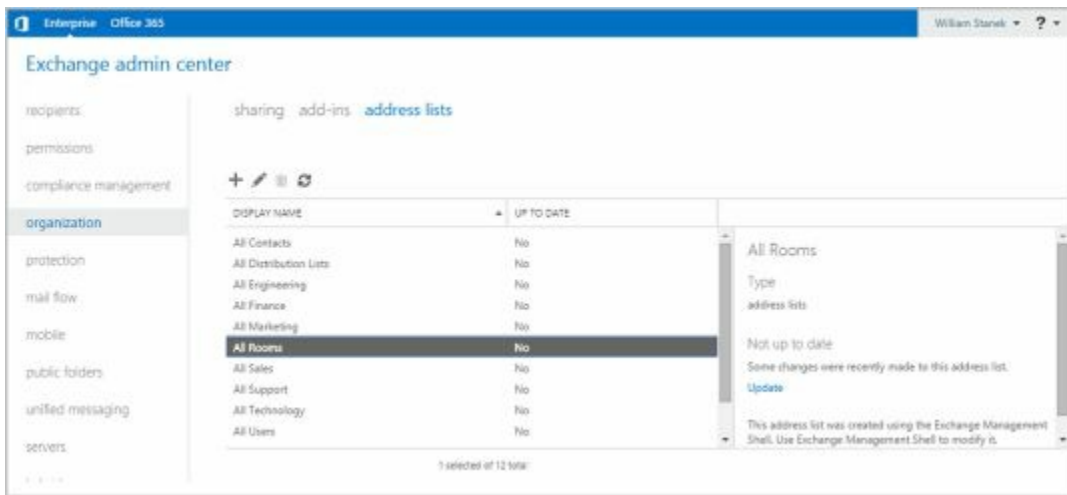


FIGURE 9-1 Accessing address lists in Exchange Admin Center.

Using Address Book Policies

Most Exchange organizations don't need address book policies. However, when multiple companies share one Exchange organization, you may want to segment the global address list to provide customized (scoped) views of recipient data to users in each separate company. You segment the global address list using address book policies. Each address book policy contains a global address list, an offline address book, a room list, and one or more address lists.

You use address book policies when you need to complete separation of the recipient data. Consider the following example:

TV Press merges with Reagent Press, resulting in a merged company called RP Media. While publicly a single company, internally the television and publishing operations are distinct and separate. The only overlap between the operations is in the top-level executive team.

The company has a single Exchange organization and wants those that work in one part of the operation to have access only to recipients and resources in that operation. Employees get scoped views of All Users, All Groups, and All Rooms as well as the default global address list and the default offline address list. These scoped views include only those that work as part of the television or publishing operations and not both.

The top-level executives and their direct support staff have access to the original, unscoped address lists. This ensures they can access recipients and resources in both operations areas.

Keep in mind that the need for custom views of recipients doesn't mean that your organization needs address book policies. You can create new address lists at any time and those address lists can be scoped however you'd like them to be scoped. For example, you could create an address list called All Marketing that includes only employees in the marketing department.

You can assign address book policy to recipients in both on-premises and online Exchange organizations. Before you can use address book policy, you must do the

following:

1. Install and enable the Address Book Policy Routing Agent using these commands:

```
Install-TransportAgent -Name "ABP Routing Agent"  
-TransportAgentFactory "Microsoft.Exchange.Transport.Agent  
.AddressBookPolicyRoutingAgent  
.AddressBookPolicyRoutingAgentFactory"  
-AssemblyPath "C:\Program Files\Microsoft\Exchange Server\V15\  
TransportRoles\Agents\AddressBookPolicyRoutingAgent\  
Microsoft.Exchange.Transport.Agent  
.AddressBookPolicyRoutingAgent.dll"
```

```
Enable-TransportAgent "ABP Routing Agent"
```

NOTE Here, C:\Program Files\Microsoft\Exchange Server\V15\ is the Exchange install path. If you installed Exchange in a different location, revise the path as appropriate.

2. Next, you should restart the Microsoft Exchange Transport service and enable Address Book Policy routing in the organization using these commands:

```
Restart-Service MExchangeTransport  
Set-TransportConfig -AddressBookPolicyRoutingEnabled $true
```

3. Set an attribute on all recipients that can be used to segment the Exchange organization. For example, you could use a custom attribute to do this.
4. Create a one or more address lists that provide the segmented views of the organization-wide global address list. Typically, you'll want a list for mailbox users, contacts, distributions lists, and rooms. Use New-AddressList with recipient filters that look for the special attribute to create these lists. Here are examples:

```
New-AddressList -Name "B - All Users" -RecipientFilter  
{(RecipientType -eq 'UserMailbox') -and  
(CustomAttribute8 -eq "CompanyB")}
```

```
New-AddressList -Name "B - All Contacts" -RecipientFilter  
{((RecipientType -eq "MailUser") -or (RecipientType  
-eq "MailContact")) -and (CustomAttribute8 -eq "CompanyB")}
```

```
New-AddressList -Name "B - All Groups" -RecipientFilter  
{((RecipientType -eq "MailUniversalDistributionGroup")  
-or (RecipientType -eq "DynamicDistributionGroup")  
-or (RecipientType -eq "MailUniversalSecurityGroup"))  
-and (CustomAttribute1 -eq "CompanyB")}
```

```
New-AddressList -Name "B - All Rooms" -RecipientFilter {(Alias  
-ne $null) -and (CustomAttribute8 -eq "CompanyB")-and  
(RecipientDisplayType -eq 'ConferenceRoomMailbox') -or
```

```
(RecipientDisplayType -eq 'SyncedConferenceRoomMailbox'})
```

NOTE Address book policy requires a room list. If you don't use rooms, create an empty list.

5. Create a segmented global address list and then use this address list to create the segmented offline address book. Here are examples:

```
New-GlobalAddressList -Name "B - GAL" -RecipientFilter  
{{CustomAttribute8 -eq "CompanyB"}}
```

```
New-OfflineAddressBook -Name "B - OAB" -AddressLists "\B - GAL"
```

6. Create an address book policy for the first company within the organization and then assign this policy to the appropriate mailboxes. Here are examples:

```
New-AddressBookPolicy -Name "CompanyB ABP" -AddressLists  
"\B - All Users", "\B - All Contacts", "\B - All Contacts"  
-OfflineAddressBook "\B - OAB" -GlobalAddressList "\B - GAL"  
-RoomList "\B - All Rooms"
```

```
Get-Mailbox -resultsize unlimited | where {$_.CustomAttribute8 -eq  
"CompanyB"} | Set-Mailbox -AddressBookPolicy "CompanyB ABP"
```

7. As necessary, repeat Steps 4 through 6 to configure address lists and policies for each company within the Exchange organization.

Creating and Applying New Address Lists

You can create new address lists to create customized views of recipient data. For example, if your organization has offices in Seattle, Portland, and San Francisco, you might want to create separate address lists for each office.

To create an address list that users can select in their Outlook clients, follow these steps:


1. In Exchange Admin Center, select **Organization** in the Navigation menu and then select **Address Lists**.
2. Click **New** (). This opens the New Address List dialog box.
3. Type an internal Exchange name and a display name for the address list, as shown in Figure 9-2. The display name should describe the types of recipients that are viewed through the list. For example, if you're creating a list for recipients in the Boston office, you can call the list Boston Office.

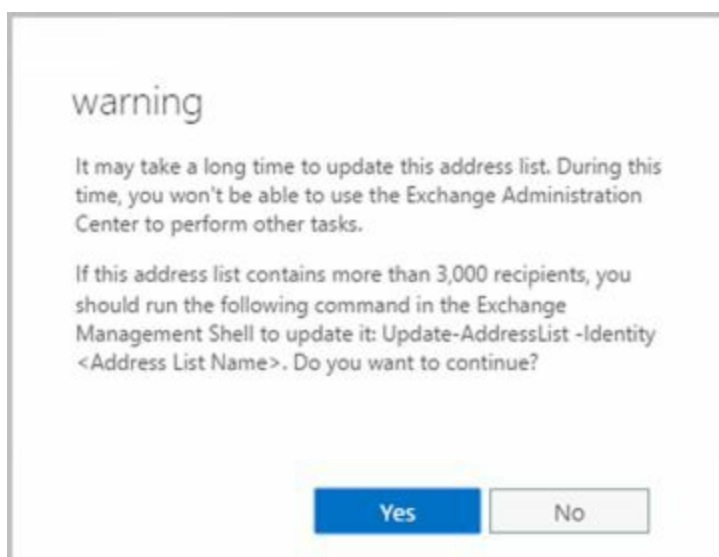
FIGURE 9-2 Specifying a name and configuring the address list.

4. The container on which you base the address list sets the scope of the list. The list will include recipients in address lists in and below the specified container. The default (root) container, \, specifies that all address lists are included by default. To specify a different container for limiting the list scope, click **Browse**, and then use the Address List Picker dialog box to select a container. In most cases, you'll want to select the default (root) container. The list path is fixed when you create a list, so you won't be able to specify a different list path later.
5. Use the Types Of Recipient To Include options to specify the types of recipients to include in the address list. Select All Recipient Types or select Only The Following Recipient Types and then select the types of recipients. You can include mailbox users, mail-enabled contacts, mail-enabled groups, mail-enabled users, and resource mailboxes.
6. Next, you can create rules that further filter the address list. Each rule acts as a condition that must be met. If you set more than one rule, each condition must be met for there to be a match. To define a rule, click **Add A Rule** and then set the filter conditions. The following types of conditions are available as well as conditions for custom attributes:
 - **Recipient Container** Filters recipients based on where in Active Directory the related account is stored. Selecting this option displays the Select An Organizational Unit dialog box. Click the container where the recipients are stored, such as Users or an organizational unit, and then click OK.
 - **State Or Province** Filters recipients based on the value of the State/Province text

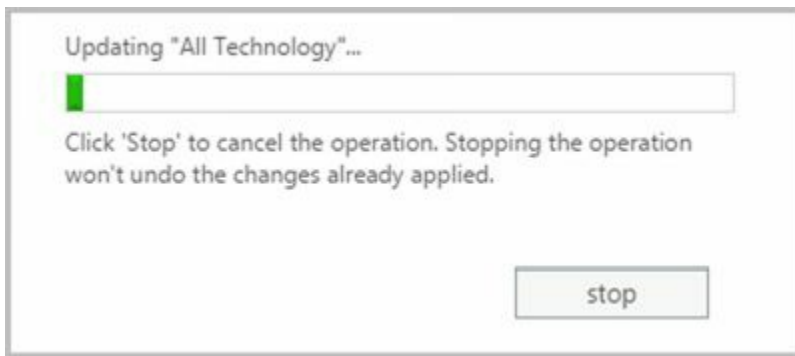
box on the Contact Information page in the related Properties dialog box. Selecting this option displays the Specify Words Or Phrases dialog box. Type a state or province identifier to use as a filter condition and then press Enter or click Add. Repeat as necessary, and then click OK.

- **Department** Filters recipients based on the value of the Department text box on the Organization page in the related Properties dialog box. Selecting this option displays the Specify Words Or Phrases dialog box. Type a department name to use as a filter condition and then press Enter or click Add. Repeat as necessary, and then click OK.
 - **Company** Filters recipients based on the value of the Company text box on the Organization page in the related Properties dialog box. Selecting this option displays the Specify Words Or Phrases dialog box. Type a company name to use as a filter condition and then press Enter or click Add. Repeat as necessary, and then click OK.
7. Click **Save** to create the address list. After the address list is created, users will be able to use the new address list the next time they start Outlook. In the Details pane, the new list will have a status of Not Up To Date.

Creating and fully populating address lists can be resource intensive, so new address lists aren't populated. You can populate the address list for the first time by updating it. To do this, click the address list and then click Update.



As shown, you'll see a warning prompt explaining that it could take a long time to update the address list. When you click Yes, Exchange Admin Center begins updating the address list and displays the update progress in a bar graph. If you find the update is taking too long, you can click Stop to halt the update. You can then restart the update process later.



In Exchange Management Shell, creating and applying address lists are two separate tasks. You can create address lists using the `New-AddressList` cmdlet. You apply address lists using the `Update-AddressList` cmdlet. Sample 9-1 provides the syntax and usage for the `New-AddressList` cmdlet. Sample 9-2 provides the syntax and usage for the `Update-AddressList` cmdlet. For `-IncludedRecipients`, you can include mailbox users, mail-enabled contacts, mail-enabled groups, mail-enabled users, and resource mailboxes.

TIP Exchange Server 2016 does not support Recipient Update Service (RUS). To replace the functionality of RUS, you can schedule the `Update-AddressList` and `Update-EmailAddressPolicy` cmdlets to run periodically using Task Scheduler. Alternatively, you can run the cmdlets manually when you modify addresses.

SAMPLE 9-1 `New-AddressList` cmdlet syntax and usage

Syntax

```
New-AddressList -Name ListName [-Container BaseAddressList]  
[-DisplayName DisplayName] [-IncludedRecipients <None, MailboxUsers,  
MailContacts, MailGroups, MailUsers, Resources, AllRecipients>]  
[-ConditionalCompany CompanyNameFilter1 , CompanyNameFilter2 ,... ]  
[-ConditionalCustomAttribute X Value1, Value2, ... ]  
[-ConditionalDepartment DeptNameFilter1 , DeptNameFilter2 , ... ]  
[-ConditionalStateOrProvince StateFilter1, StateFilter2, ... ]  
[-DomainController FullyQualifiedName ] [-Organization OrgName ]  
[-RecipientContainer ApplyFilterContainer ]
```

```
New-AddressList -Name ListName [-Container BaseAddressList]  
[-DisplayName DisplayName ] [-DomainController FullyQualifiedName ]  
[-Organization OrgName ] [-RecipientContainer ApplyFilterContainer ]  
[-RecipientFilter Filter ]
```

Usage

```
New-AddressList -Name 'West Coast Sales' -Container '\'  
-DisplayName 'West Coast Sales' -IncludedRecipients 'MailboxUsers,  
MailContacts, MailGroups, Resources'  
-ConditionalCompany 'Imagined Lands'  
-ConditionalDepartment 'Sales','Marketing'  
-ConditionalStateOrProvince 'Washington','Idaho','Oregon'
```

SAMPLE 9-2 Update-AddressList cmdlet syntax and usage

Syntax

```
Update-AddressList -identity ListIdentity  
[-DomainController FullyQualifiedName ]
```

Usage

```
Update-AddressList -Identity '\\West Coast Sales'
```

Updating Address List Configuration and Membership Throughout the Domain


Exchange Server doesn't immediately replicate changes to address lists throughout the domain. Instead, changes are replicated during the normal replication cycle, which means that some servers might temporarily have outdated address list information. Rather than waiting for replication, you can manually update address list configuration, availability, and membership throughout the domain. To do this, follow these steps:

1. In Exchange Admin Center, select **Organization** in the Navigation menu and then select **Address Lists** .
2. Click the address list you want to work with and then click **Update** .
3. You'll see a warning prompt explaining that it could take a long time to update the address list. Click **Yes** . Exchange Admin Center begins updating the address list and displays the update progress in a bar graph.
4. If you find the update is taking too long, you can click **Stop** to halt the update. You can then restart the update process later.

Alternatively, you can use the Update-AddressList cmdlet to update lists. See Sample 9-2 for syntax and usage.

Previewing and Editing Address Lists

Although you can't change the properties of default address lists, you can change the properties of address lists that you create using either Exchange Admin Center or Exchange Management Shell. You can edit a list's settings or preview the recipients in the list by completing the following steps:

1. In Exchange Admin Center, select **Organization** in the Navigation menu and then select **Address Lists** .
2. Click the address list you want to work with. If there's a note in the Details pane stating the list was created in Exchange Management Shell, you won't be able to modify its settings. You can, however, view the list's settings in the Address List dialog box.
3. Click **Edit** () . In the Address List dialog box, you'll see the name, path, and recipient filter associated with the list.

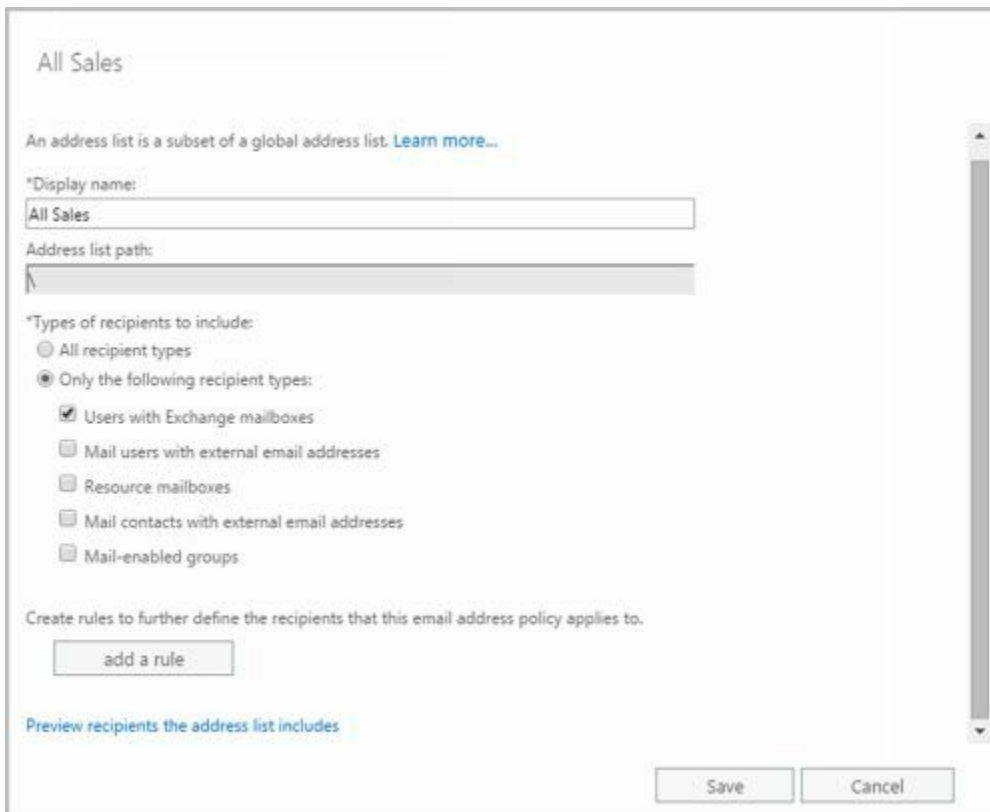
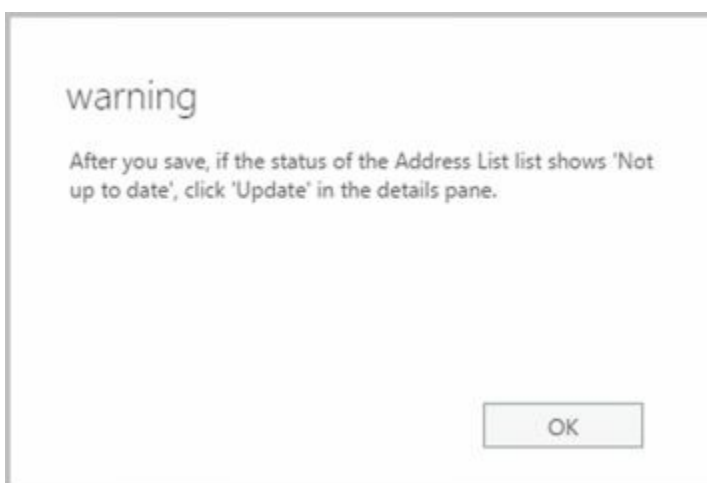


FIGURE 9-3 Modifying an address list.

4. To preview the recipients included in the list, click the link provided.
5. Modify the name as necessary. Use the Types Of Recipients To Include options to specify the types of recipients to include. Select **All Recipient Types** or select **Only The Following Recipient Types** and then select the types of recipients.
6. Create new rules or modify existing rules to further filter the recipients.
7. Click **Save** . A warning prompt is displayed stating the address list might not be up to date if you made changes. Click **OK**.



In Exchange Management Shell, you can modify an address list using the Set-AddressList cmdlet. Sample 9-3 provides the syntax and usage. When you modify an address list, you can make the changes visible by using the Update-AddressList cmdlet, as shown previously in Sample 9-2.

SAMPLE 9-3 Set-AddressList cmdlet syntax and usage

Syntax

```
Set-AddressList -Identity ListName
[-DisplayName DisplayName ] [-IncludedRecipients <None, MailboxUsers,
MailContacts, MailGroups, Resources, AllRecipients> ]
[-ConditionalCompany CompanyNameFilter1 , CompanyNameFilter2 ,... ]
[-ConditionalDepartment DeptNameFilter1 , DeptNameFilter2 , ... ]
[-ConditionalStateOrProvince StateFilter1 , StateFilter2 , ... ]
[-DomainController FullyQualifiedName ] [-ForceUpgrade <$false|$true>]
[-RecipientContainer ApplyFilterContainer ] [-RecipientFilter Filter ]
```

Usage

```
Set-AddressList -Identity '\\West Coast Sales' -Name 'Sales Team-West'
-IncludedRecipients 'MailboxUsers, MailContacts, MailGroups'
-Company 'Imagined Lands'
-Department 'Sales','Marketing'
-StateOrProvince 'Washington','Idaho','Oregon'
```

Usage

```
Set-AddressList -Identity '\\West Coast Sales' -Name 'Sales Team-West'
-IncludedRecipients 'MailboxUsers, MailContacts, MailGroups'
-ForceUpgrade $true
```

Configuring Clients to Use Address Lists

Address books are available to clients that are configured for corporate or workgroup use. To set the address lists used by the client, complete these steps:



1. In Office Outlook 2013 or Outlook 2016, on the Home panel, select **Address Book** . Alternatively, press **Ctrl+Shift+B** .
2. In the Address Book dialog box, from the Tools menu, select **Options** , and then set the following options to configure how address lists are used:
 - **When Sending E-Mail, Check Address Lists In ThisOrder** Sets the order in which Outlook searches address books when you send a message or click Check Names. You can start with either the global address list or the contact folders. Or you can choose the Custom option and then use the up and down arrows to change the list order.
 - **When Opening The Address Book, Show This Address ListFirst** Sets the address book that the user sees first whenever he or she works with the address book.
3. Click **OK** .

TIP When checking names, you'll usually want the global address list (GAL) to be listed before the user's own contacts or other types of address lists. This is important because users often put internal mailboxes in their personal address lists. The danger of doing this without first resolving names against the GAL is that

although the display name might be identical, the properties of a mailbox might change. When changes occur, the entry in the user's address book is no longer valid, and any mail sent bounces back to the sender with a nondelivery report (NDR). To correct this, the user should either remove that mailbox from his or her personal address list and add it based on the current entry in the GAL, or change the check names resolution order to use the GAL before any personal lists.

Renaming and Deleting Address Lists

You can only rename or delete user-defined address lists.

- **Renaming address lists** To rename an address list, in Exchange Admin Center, select its entry and then select **Edit** (). Type a new name in the Display Name text box. In the Details pane, the modified list will have a status of Not Up To Date. To update the membership of the address list, click **Update**.
- **Deleting address lists** To delete an address list, in Exchange Admin Center, select its entry and then select **Remove** (). When prompted to confirm the action, click **Yes**.

In Exchange Management Shell, you can remove address lists using the Remove-AddressList cmdlet. Sample 9-4 provides the syntax and usage. If you also want to remove address lists that reference the address list you are removing and match a portion of it (child address lists), set the -Recursive parameter to \$true. By default, the cmdlet does not remove child address lists of the specified list.

SAMPLE 9-4 Remove-AddressList cmdlet syntax and usage

Syntax

```
Remove-AddressList -Identity ListIdentity  
[-DomainController FullyQualifiedName] [-Recursive {$true | $false}]
```

Usage

```
Remove-AddressList -Identity '\West Coast Sales'
```

Managing Offline Address Books

You configure offline address books differently than online address lists. To use an offline address book, the client must be configured to have a local copy of the server mailbox, or you can use personal folders. Clients using Outlook 2010 or later retrieve the offline address book from the designated offline address book (OAB) distribution point.

NOTE Although future updates may change this, Exchange Admin Center doesn't have options for managing offline address books at the time of this writing. This means that you need to use Exchange Management Shell to manage offline address books.

IMPORTANT An OAB distribution point is a virtual directory to which Outlook 2010 or later clients can connect to download the offline address book. OAB distribution points are hosted by Mailbox servers running Internet Information Services (IIS) as virtual directories. Each distribution point can have two URLs associated with it: one URL for internal (on-site) access and another for external (off-site) access.

Creating Offline Address Books

By default, the default offline address book includes all the addresses in the global address list. It does this by including the default global address list. All other offline address books are created by including the default global address list or a specific online address list as well.

NOTE You can create other custom offline address books using Exchange Management Shell. You cannot use Exchange Admin Center to create other offline address books.

In Exchange Management Shell, you can create offline address books using the `New-OfflineAddressBook` cmdlet. You apply offline address books using the `Update-OfflineAddressBook` cmdlet. Sample 9-5 provides the syntax and usage for the `New-OfflineAddressBook` cmdlet. Sample 9-6 provides the syntax and usage for the `Update-OfflineAddressBook` cmdlet.

NOTE Public folder distribution is no longer associated with offline address books. Public folders are now stored in special mailboxes, as discussed in Chapter 6, "Adding Special-Purpose Mailboxes."

SAMPLE 9-5 `New-OfflineAddressBook` cmdlet syntax and usage

Syntax

```
New-OfflineAddressBook -Name ListName -Server GenerationServer
```


-AddressLists **AddressList1** , **AddressList2** , ...
[-VirtualDirectories **VirtualDir1** , **VirtualDir2** , ...] {AddtlParams}

{AddtlParams}
[-DiffRetentionPeriod **RetentionPeriod**]
[-DomainController **FullyQualifiedName**]
[-GlobalWebDistributionEnabled <\$true | \$false>]
[-IsDefault <\$true | \$false>] [-Organization **OrgName**]

Usage

```
New-OfflineAddressBook -Name ' Offline – West Coast Sales '  
-Server ' CorpSvr127 '  
-AddressLists ' \West Coast Sales '  
-VirtualDirectories 'CORPSVR127\OAB (Default Web Site)'
```

SAMPLE 9-6 Update-OfflineAddressBook cmdlet syntax and usage

Syntax

```
Update-OfflineAddressBook-Identity OABName  
[-DomainController FullyQualifiedName ]
```

Usage

```
Update-OfflineAddressBook -Identity '\Offline – West Coast Sales'
```

When you create an offline address book, you must use the -AddressLists parameter to specify the address lists that are included. If you want the offline address book to include all recipients in the organization, specify that the Default Global Address List is the address list to include as shown in this example:

```
New-OfflineAddressBook -Name ' Offline – Entire Organization '  
-Server ' CorpSvr127 '  
-AddressLists '\Default Global Address List'  
-VirtualDirectories 'CORPSVR127\OAB (Default Web Site)'
```

You can include multiple address lists using a comma-separated list, as shown in this example:

```
New-OfflineAddressBook -Name ' Offline – Sales & Marketing '  
-Server ' CorpSvr127 '  
-AddressLists '\All Marketing', '\All Sales', '\Sales Teams'  
-VirtualDirectories 'CORPSVR127\OAB (Default Web Site)'
```

If you want the new offline address book to be the default, use the -IsDefault parameter.

Configuring Clients to Use an Offline Address Book

Offline address lists are available only when users are working offline. You can configure how clients use offline address books by completing the following steps:

1. Do one of the following:

- In Outlook 2010, click the Office button. On the Info pane, select Download Address Book. The Offline Address Book dialog box appears.
- In Outlook 2013 and Outlook 2016, on the File pane, click Info. On the Info page, click Account Settings and then select Download Address Book. The Offline Address Book dialog box appears.
 2. Select the Download Changes Since Last Send/Receive check box to download only items that have changed since the last time you synchronized the address list. Clear this check box to download the entire contents of your address book.
 3. Specify the information to download as either of the following two options:
 - **Full Details** Select this option to download the address book with all address information details. Full details are necessary if the user needs to encrypt messages when using remote mail.
 - **No Details** Select this option to download the address book without address information details. This reduces the download time for the address book.
 4. If multiple address books are available, use the **Choose Address Book** drop-down list to specify which address book to download. Click **OK**.

Setting the Default Offline Address Book

Although you can create many offline address books, clients download only one. This address list is called the default offline address book. To specify the default offline address book, use Set-OfflineAddressBook with this basic syntax:

```
Set-OfflineAddressBook -Identity OABName -IsDefault
[-DomainController FullyQualifiedName ]
```

In the following example, Offline – All Company is set as the default offline address book:

```
Set-OfflineAddressBook -Identity '\Offline – All Company' -IsDefault
```

Changing Offline Address Book Properties

The offline address book is based on other address lists that you've created in the organization. In Exchange Management Shell, you can modify offline address books using the Set-OfflineAddressBook cmdlet. Sample 9-7 provides the syntax and usage.

SAMPLE 9-7 Set-OfflineAddressBook cmdlet syntax and usage

Syntax

```
Set-OfflineAddressBook -Identity OABName
[-AddressLists AddressList1 , AddressList2 , ... ]
[-ApplyMandatoryProperties {$true | $false}]
[-ConfiguredAttributes Attributes ]
[-DiffRetentionPeriod RetentionPeriod ]
[-DomainController FullyQualifiedName ]
[-GlobalWebDistributionEnabled <$true | $false>]
```

```
[-IsDefault <$true | $false>] [-MaxBinaryPropertySize Size ]
[-MaxMultivaluedBinaryPropertySize Size ]
[-MaxMultivaluedStringPropertySize Size ] [-MaxStringPropertySize Size ]
[-Name Name ] [-PublicFolderDistributionEnabled <$false|$true> ]
[-Schedule Schedule ] [-UseDefaultAttributes {$true | $false}]
[-Versions Versions ] [-VirtualDirectories VirtualDir1 , VirtualDir2 , ...]
```

Usage

```
Set-OfflineAddressBook -Identity '\Offline – West Coast Sales'
-Name 'West Coast Sales - Offline'
-AddressLists '\West Coast Sales'
-PublicFolderDistributionEnabled $true
-VirtualDirectories 'CORPSVR127\OAB (Default Web Site)'
```

One way to modify an offline address book is to modify the list of included address lists. You can make additional address lists a part of the offline address book. If you no longer want an address list to be a part of the offline address book, you can remove it. To perform either task, use the `-AddressLists` parameter. This parameter specifies the exact list of address lists to include, and you must always explicitly specify each address list that should be included. Consider the following example:

```
Get-OfflineAddressBook
```

Name	Versions	AddressLists
-----	-----	-----
Default Offline Address Book	{Version4}	{\Default Global Address List}
Temp Employees Address Book	{Version4}	{\All Support, \All Temps}

In this example, the organization has two offline address books. One for full-time employees and one for temporary employees who provide onsite support. For temporary employees, the offline address book includes recipient data only for members of the support team and other temps on the support team. If the offline address book for temporary employees should also include recipient data for All Help Desk, you could add this address list as shown in this example:

```
Set-OfflineAddressBook -Identity '\Temp Employees Address Book'
-AddressLists '\All Support', '\All Temps', 'All Help Desk'
```

If you later decided to remove All Help Desk from this offline address book, you could do so by entering the following command:

```
Set-OfflineAddressBook -Identity '\Temp Employees Address Book'
-AddressLists '\All Support', '\All Temps'
```

Designating OAB Generation Servers and Schedules

In Exchange 2016, the organization has a dedicated OAB generation server. This server is responsible for generating the offline address books for the entire organization. Although the first Mailbox server you install with Exchange 2016 may be designated as the OAB generation server, this isn't always the case.

To identify the OAB generation server, you need to locate the arbitration mailbox that handles the offline address book generation. In Exchange 2016, an arbitration mailbox with the persisted capability "OrganizationCapabilityOABGen" handles offline address book generation. You can locate this mailbox and identify the server and database it resides on using the following command:

```
Get-Mailbox -Arbitration | where {$_.PersistedCapabilities -like "*oab*"} |  
ft name, servername, database
```

If your Mailbox servers are configured in an availability group, ensure you've identified the active copy of the database using the following command:

```
Get-MailboxDatabaseCopyStatus DatabaseName
```

where DatabaseName is the database to check. The active copy has the status Mounted.

By default, the OAB generation server rebuilds offline address books on a daily schedule and does so once each day. You can confirm the current settings using the following command:

```
Get-MailboxServer -Identity OABGenerationServer | fl OABGeneratorWorkCycle,  
OABGeneratorWorkCycleCheckpoint
```

where OABGenerationServer is the Mailbox server hosting the OAB generation mailbox. The output of this command is as shown here:

```
OABGeneratorWorkCycle          : 1.00:00:00  
OABGeneratorWorkCycleCheckpoint : 1.00:00:00
```

NOTE In Exchange 2010, offline address book generation occurred according to a fixed schedule set with the -Schedule parameter of Set-OfflineAddressBook. In Exchange 2016, this schedule is not used.

The Mailbox server uses the default daily schedule and will rebuild the offline address books once each day. The schedule uses the following format:

D.HH:MM:SS

where D is the number of days, HH sets the hours, MM sets the minutes, and SS sets the seconds.

You can configure a different schedule using Set-MailboxServer. Use -OABGeneratorWorkCycle to set the master schedule and -OABGeneratorWorkCycleCheckpoint to set the rebuild interval within this schedule. For example, if you want address books to be rebuild daily and update every six hours, use the following command:

```
Set-MailboxServer -OABGeneratorWorkCycle 1.00:00:00  
-OABGeneratorWorkCycleCheckpoint 06:00:00
```

The OAB generation server manages and propagates the offline address books. If the OAB generation server is being overutilized and you want to move the offline address book generation responsibility to a server with more resources, you can do this using

several different techniques. When the database and server are part of an availability group, you can move the OAB generation mailbox from one server in the group to another server in the group. However, to do this, you must activate the corresponding mailbox database on the other server (and thereby inactivate the mailbox database on its current server). Consider the following example:

```
Move-ActiveMailboxDatabase Database42 -ActivateOnServer MailServer22
```

In this example, MailServer22 hosts an inactive copy of the mailbox database that contains the OAB generation mailbox and this database is activated. When Database42 is activated, MailServer22 becomes the OAB generation server.

When the database and server are not part of an availability group, you can use a standard move request to move the OAB generation mailbox from a database on one server to a database on another server. Consider the following example:

```
Get-Mailbox -Arbitration -database Database42 | where  
{$_PersistedCapabilities -like "*oab*"} | New-MoveRequest  
-TargetDatabase Database14
```

When the move request is completed and final, the new server becomes the OAB generation server. As may be required for load balancing, fault tolerance, or geographically disbursed Exchange organizations, you can create an additional OAB generation mailbox. To do this, use the following commands:

```
New-Mailbox -Arbitration -Name "OAB 2" -Database Database42  
-UserPrincipalName oab2@imaginedlands.com -DisplayName "OAB Mailbox 2"
```

```
Set-Mailbox -Arbitration oab2 -OABGen $true
```

Rebuilding the OAB

Although the offline address book is generated automatically according to the generator work cycle, you can force the OAB generator to rebuild offline address books manually. To do this, use the Update-OfflineAddressBook cmdlet as shown in this example:

```
Update-OfflineAddressBook -Identity '\Default Offline Address Book'
```

This example initiates an update of the default offline address book. This command initiates an RPC request to each mailbox server hosting an active OAB generation mailbox.

You also can force Exchange to rebuild the offline address book if you restart the Mailbox Assistance service on the server hosting an active OAB generation mailbox.

Deleting Offline Address Books

If an offline address book is no longer needed, you can delete it as long as it isn't the default offline address book. Before you can delete the default offline address book, you must set another address book as the default.

In Exchange Management Shell, you can delete an offline address book using the `Remove-OfflineAddressBook` cmdlet. Sample 9-8 provides the syntax and usage. Set the `-Force` parameter to `$true` to force the immediate removal of an offline address book.

SAMPLE 9-8 `Remove-OfflineAddressBook` cmdlet syntax and usage

Syntax

```
Remove-OfflineAddressBook -Identity ' OfflineAddressBookIdentity'  
[-Force <$false|$true>] [-DomainController FullyQualifiedName ]
```

Usage

```
Remove-OfflineAddressBook -Identity '\Offline – West Coast Sales'
```

Chapter 10. Configuring Exchange Clients

Knowing how to configure and maintain Exchange clients is essential for Microsoft Exchange administrators. With Microsoft Exchange Server 2016 and Exchange Online, you can use any mail client that supports standard mail protocols. For ease of administration, however, you'll want to choose specific clients for users. I recommend focusing on Microsoft Office Outlook 2010 and later and Outlook Web App as your clients of choice. Each client supports a slightly different set of features and messaging protocols, and each client has its advantages and disadvantages, including the following:

- With Outlook 2010 or later, you get a full-featured client that on-site, off-site, and mobile users can use. Outlook 2010 or later is part of the Microsoft Office system of applications. They are the only mail clients that support the latest messaging features in Exchange Server. Corporate and workgroup users often need their rich support for calendars, scheduling, voice mail, and email management.
- With Outlook Web App, you get a mail client that you can access securely through a standard web browser whether you are using Windows desktop, Windows Server, iOS or Android. With Microsoft Edge, Internet Explorer 11.0 or later, and current versions of Firefox, Chrome and Safari, Outlook Web App supports many of the features found in Outlook 2010 and later, including calendars, scheduling, and voice mail. With other browsers, the client functionality remains the same, but some features might not be supported. You don't need to configure Outlook Web App on the client, and it's ideal for users who want to access email while away from the office.

Outlook 2010 and later versions are the most common Exchange clients for corporate and workgroup environments. With the MAPI over HTTP feature of Exchange, which eliminates the need for a virtual private network (VPN) to securely access Exchange Server over the Internet by using the Messaging Application Programming Interface (MAPI) over Secure Hypertext Transfer Protocol (HTTPS) connections, Outlook 2010 and later versions might also be your clients of choice for off-site and mobile users.

NOTE With Exchange 2016 and Exchange Online, MAPI over HTTP is enabled by default in a standard configuration. However, if you upgraded to Exchange 2016 from Exchange 2013, you'll need to enable MAPI over HTTP, as the feature was disabled by default in Exchange 2013. Note also that only Outlook 2016 supports MAPI over HTTP without any updates. Outlook 2010 requires Service Pack 2 and updates KB2956191 and KB2965295 to support MAPI over HTTP. Outlook 2013 requires Service Pack 1.

This chapter shows you how to manage Outlook 2010 and later. For ease of reference, I will refer to Outlook 2010 and later simply as Outlook, unless I need to differentiate between them.

Mastering Outlook Web App essentials

Outlook Web App is a standard Microsoft Exchange Server 2016 technology that allows users to access their mailboxes using a web browser. If public folders are hosted by Exchange 2016, users will be able to access public folder data as well. The technology works with standard Internet protocols, including HTTP and Secure HTTP (HTTPS).

When users access mailboxes and public folder data over the web, Client Access and Mailbox servers are working behind the scenes to grant access and transfer files to the browser. Because you don't need to configure Outlook Web App on the client, it's ideally suited for users who want to access email while away from the office and may also be a good choice for users on the internal network who don't need the full version of Microsoft Outlook. Outlook Web App is automatically configured for use when you install the Client Access and Mailbox server roles for Exchange Server 2016. This makes Outlook Web App easy to manage. That said, there are some essential concepts you should know to manage Outlook Web App more effectively, and the following section explains these concepts.

Getting started with Outlook Web App

Outlook Web App is installed automatically when you install the Mailbox server role for Exchange Server 2016. If users will be accessing Outlook Web App over the Internet, the server must be able to accept connections from external clients on an external URL.

In most cases, you need to open only TCP port 443 on your organization's firewall to allow users to access mailboxes and public folder data over the web. After that, you simply tell users the URL path that they need to type into their browser's Address text box in order to access Outlook Web App when they're off-site.

Outlook Web App for Exchange 2016 has a streamlined interface that is optimized for PCs, tablets, and mobile devices. The browser used to access Outlook Web App determines the experience and supported features. The following two versions are available:

- **Standard** Provides a rich experience with performance that closely approximates Microsoft Office Outlook, including a folder hierarchy that you can expand or collapse, drag-and-drop functionality, move and copy functionality, and shortcut menus that you can access by right-clicking. In addition, you can use all of the following features: appearance color schemes, calendar views, file share integration, notifications, personal distribution lists, public folder access, recover deleted items, reminders, search, server-side rules, voice mail options, and WebReady Document Viewing.
- **Light** Provides a basic experience with a simplified user interface when the user's browser cannot support the standard version. This version can also be useful when working over low-bandwidth connections or when there are accessibility needs. No

Standard-only features are available. In addition, calendar options are limited and messages can be composed only as plain text. Outlook Web App shortcut menus are not displayed when you right-click. The Outlook Web App toolbar has slightly different options, and the Options page itself is simplified as well.

IMPORTANT By default, all users see the standard version when their browser supports it. Additionally, Outlook Web App for Exchange 2016 doesn't include a spellchecker as this functionality is now being built into web browsers. Microsoft Internet Explorer 10 and later as well as some other web browsers have built-in spell checkers.

Outlook Web App uses HTML 4.0 and JavaScript [European Computer Manufacturers Association (ECMA)] script. The standard version of Outlook Web App is available for PCs, servers, tablets and smart phones running current versions of Windows, Android and iOS

Outlook Web App for Exchange Server 2016 has many features, including:

- **Apps** Users and administrators can add apps to the interface to add functionality. Several apps are installed and made available to users by default, including the following apps created by Microsoft: Action Items, Bing Maps, Suggested Meetings and Unsubscribe. Other apps can be added from the Office Store, from a URL, or from a file.
- **Inbox Rules** Users can create Inbox rules to automatically sort incoming email into folders. Users create rules on the Inbox Rules tab or by right-clicking a message on which they want to base a rule, and then selecting Create Rule.
- **Text Messaging Notifications** Users can set up text messaging notifications to be sent to their mobile devices. Notifications are triggered by calendar events, such as meetings and Inbox rules.
- **Message attachments** Users can attach files, meeting requests, and other messages to messages by clicking the attach file icon on the toolbar.
- **Delivery Reports** Users can generate delivery reports to search for delivery information about message they've sent or received during the previous two weeks.
- **Personal Groups** Users can create personal groups that will appear in their address book.
- **Public groups** Users can create distribution groups that will appear in the global address book for everyone to use.

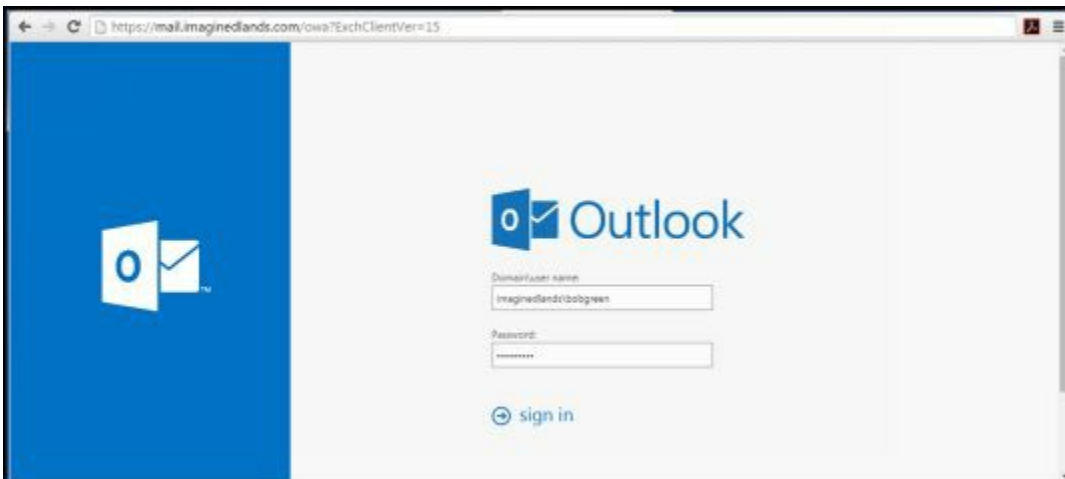
Using message options, you can specify message sensitivity as Normal, Personal, Private or Confidential. You also can request delivery receipt, read receipt or both.

Connecting to Mailboxes and Public Folder Data Over the Web

With Outlook Web App, you can easily access mailboxes and public folder data over the web and a corporate intranet. To access a user's mailbox, type the Exchange Outlook

Web App URL into your browser's Address text box, and then enter the user name and password for the mailbox you want to access. The complete step-by-step procedure is as follows:

1. In a web browser, enter the secure URL for Outlook Web App. If you are outside the corporate network, enter the external URL, such as `https://servername.yourdomain.com/owa`, where `servername` is a placeholder for the web server hosted by Exchange Server and `yourdomain.com` is a placeholder for the external domain. For example, if your Mailbox server is configured to use mail as the external DNS name and your external domain is `imaginedlands.com`, you type `https://mail.imagedlands.com/owa`.



The version of Outlook Web App displayed depends on the version of Exchange running on the Mailbox server hosting your personal mailbox. Exchange 2010 runs version 14 and you can specify this version explicitly by appending ?

ExchClientVer=14 to the internal or external URL.

Exchange 2016 runs version 15 and you can specify this version explicitly by appending **?ExchClientVer=15** to the internal or external URL. For example, if your external URL is `https://mail.imagedlands.com`, you could enter `https://mail.imagedlands.com/owa?ExchClientVer=15` as the URL.

NOTE By default, you must use HTTPS to connect. If you don't, you'll see an error stating "Access is denied." Using HTTPS ensures data transmitted between the client browser and the server is encrypted and in this way secured.

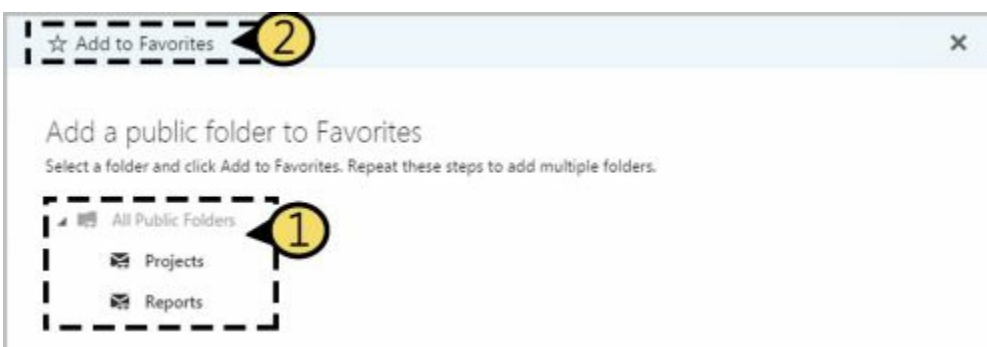
2. By default, Mailbox servers are configured to use Secure HTTP (HTTPS) for Outlook Web App. When you install Exchange Server 2016, a self-signed security certificate is issued for the Mailbox server automatically. Because this default certificate is not issued by a trusted certificate authority, you might see a warning that there is a problem with the website's security certificate. If your browser displays a security alert stating there's a problem with the site's security certificate or that the connection is untrusted, proceed anyway.
- With Internet Explorer, the error states "There's a problem with this website's security certificate". You proceed by selecting the Continue To This Web Site (Not Recommended) link.

- With Google Chrome, the error states "The site's security certificate is not trusted". You continue by selecting the Proceed Anyway button.
 - With Mozilla Firefox, the error states "This connection is untrusted." You proceed by selecting I Understand The Risks and then selecting Add Exception. Finally, in the Add Security Exception dialog box, you select Confirm Security Exception.
3. You'll see the logon page for Outlook Web App. Enter your user name and password, and then click **Sign In**.
Be sure to specify your user name in DOMAIN\username format. The domain can either be the DNS domain, such as imaginedlands.com, or the NetBIOS domain name, such as pocket-consulta. For example, the user MikeL could specify his logon name as imaginedlands.com\mikel or imaginedlands\mikel.
 4. If you are logging in for the first time, select your preferred display language and time zone, and then click **Save**.



After a user has accessed his mailbox in Outlook Web App, he can access public folders data that is available as well as long as the public folders are hosted on Exchange 2016. To access public folders, follow these steps:

1. In the left pane of the Outlook Web App window, right-click Favorites.
2. Select Add Public Folder To Favorites. In the Add Public Folder dialog box, you'll see a list of the available top levels to which you have access
3. Select a public folder and then click Add To Favorites.
4. Repeat Steps 1 through 3 to add other public folders.



The public folders you've added are listed under the Favorites heading in the left pane. To access a folder and display its contents in the main pane, simply select it in the left

pane.

Working with Outlook Web App

After you enter the Outlook Web App URL into a browser's Address text box and log in, you'll see the view of Outlook Web App compatible with your browser. Figure 10-1 shows the full-featured view of Outlook Web App.

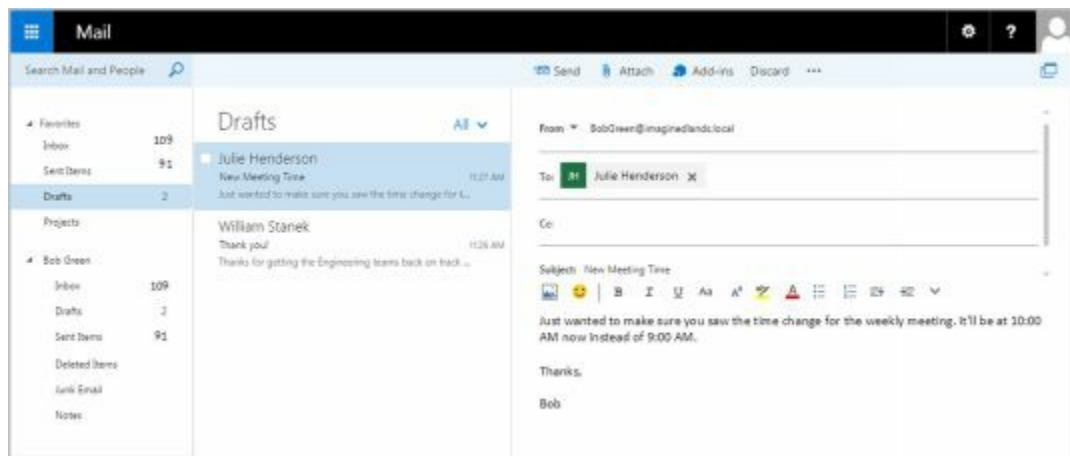



FIGURE 10-1 Outlook Web App has nearly all of the features of Microsoft Office Outlook.

Most users see this view of Outlook Web App automatically. If their browsers don't support a necessary technology for the full-featured view, some features or options won't be available, or they might see the Light view instead. If they can right-click and see a shortcut menu, they have the full-featured view.


As shown in Figure 10-1, the latest version of Outlook Web App has a toolbar that provides quick access to the following key features:



Apps – Displays a list of the available apps you can switch to, including Mail, Calendar, People and Tasks.



Help – Use this option to access online help for Outlook Web App. You can search for topics, print help text and more.



Settings – Provides quick access to settings for managing automatic replies, display settings, Outlook apps, offline settings, themes, and the user's password. Also allows the user to access the Options page to configure Outlook Web App properties or view current configuration details.



Account – Displays the user's name. Provides options for opening another mailbox and signing out. Also allows you to set the mailbox picture.



Open Window– Opens the message or other item you are working with in a separate window.

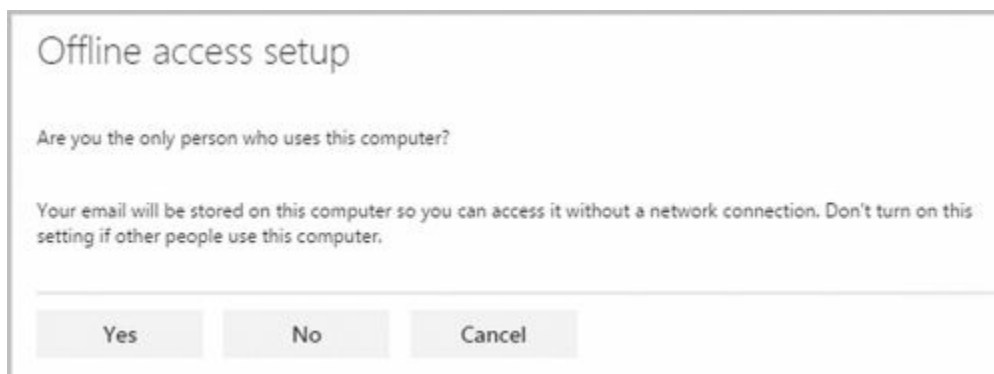
Outlook Web App can be configured to allow users to connect their account other email accounts. This allows users to keep send, receive, and read email from other email services. Users also can forward email from their Outlook Web App to another account. If users want to add their contacts from Facebook and LinkedIn to Outlook Web App contacts, Outlook Web App can be configured to do this, too.

Outlook Web App can be configured to allow users to work offline. Users can continue to work when they are disconnected from the Internet when Outlook Web App is configured to cache mail items and other information on the users' computers. When Offline mode is allowed in the Outlook Web App configuration, users can enable offline settings by completing the following steps:

1. In Outlook Web App, select Settings () , Offline Settings, choose Turn On Offline Access. This starts the Offline Settings Wizard.



2. As the cached mail and other information stored on a user's computer could be accessed by other users of a computer, the wizard prompts to ensure the current user is the only person who uses the computer. Click **Yes** to confirm.




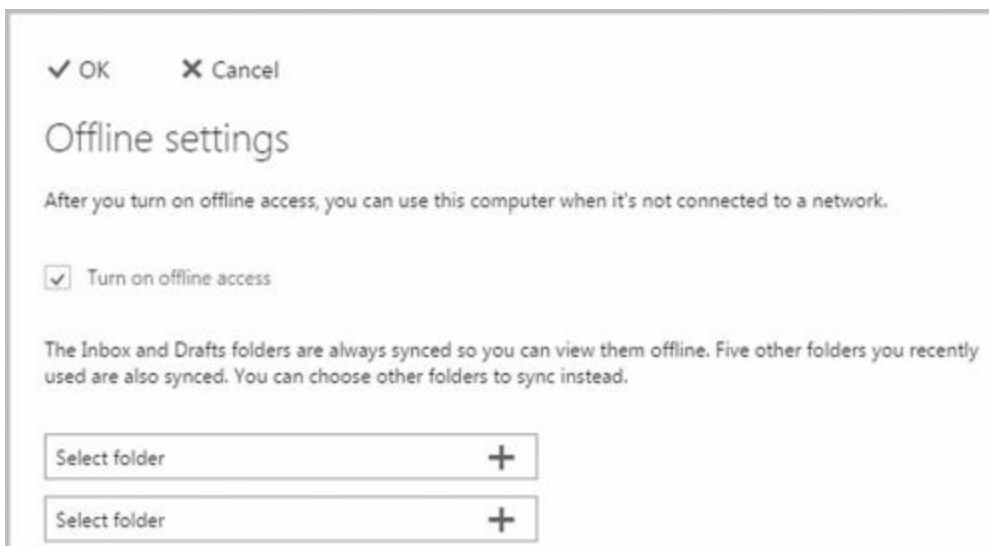
3. As a user's browser caches the mail data, the size of the browser cache and other related settings might need to be changed. If prompted to grant more storage to the

browser, click **Yes**.

4. Click **Next** to continue. When prompted, press Ctrl+D to create a bookmark for quickly accessing Outlook Web App.
5. Click **Next** and then click **OK**.

By default, the Inbox and Drafts folder, as well as recently used folders are synced for offline use. To designate folders that should always be synced:

1. In Outlook Web App, select Settings () , Offline Settings.
2. Up to five folders can be selected for syncing. Any currently selected folder is listed by name in one of the five designated slots.



You can now add or remove synced folders:

- To add a folder, click **Select Folder**. In the Select Folder dialog box, click the folder to sync, such as Sent Items, and then click **OK**.
- To remove a folder, click the Remove button () to the right of the folder name.

The primary offline data for Outlook Web App and the user's mailbox is cached under %LocalAppData%\Microsoft\Windows\WebCache on the computer. After offline access is enabled, the browser reads data from this cache, allowing users to continue to work with Outlook Web App and access mail, contacts, and other mail data when their computers aren't connected to the Internet.

If offline mode has been enabled, you can turn this feature off by:

3. In Outlook Web App, select Settings () , Offline Settings.
4. Clear the Turn On Offline Access checkbox and then click **OK**.

Disabling offline access doesn't remove the cached data, nor does clearing the browser cache. Because the cached mail data is persistent across browser sessions and

independent of the browser's local cache, you must manually remove this data if you want to be certain the data can no longer be accessed.

Enabling and Disabling Web Access for Users

Exchange Server 2016 enables Outlook Web App for each user by default and applies the Default Outlook Web App Mailbox policy to each user. Outlook Web App Mailbox policy controls the features that are enabled for each user and allows users to:

- Use Instant Messaging, text messaging, unified messaging, and Exchange Active Sync
- Create and manage personal contacts, and access all internal address lists
- Use Journaling, notes, Inbox rules, and recover deleted items
- Change their password and configure junk email filters
- Use themes, the premium client, and email signatures
- Manage calendars, tasks, reminders, and notifications

If necessary, you can enable or disable Outlook Web App or set a new default policy for specific users by completing the following steps:

1. In Exchange Admin Center, select Recipients in the Feature pane, and then select Mailboxes. You should now see a list of users with Exchange mailboxes in the organization.
2. Select the user you want to work with in the main pane.
3. In the details pane, the current status of Outlook Web App is listed under the Email Connectivity heading, as shown in Figure 10-2.

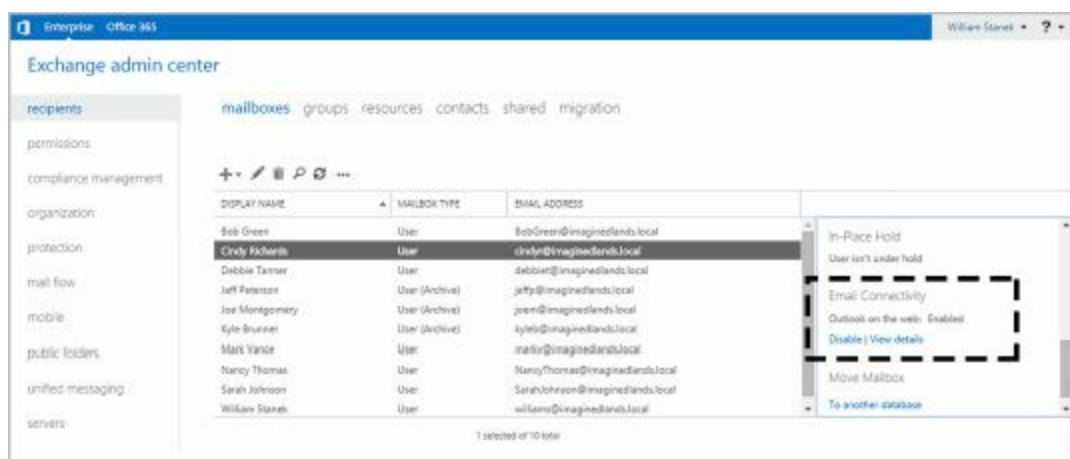
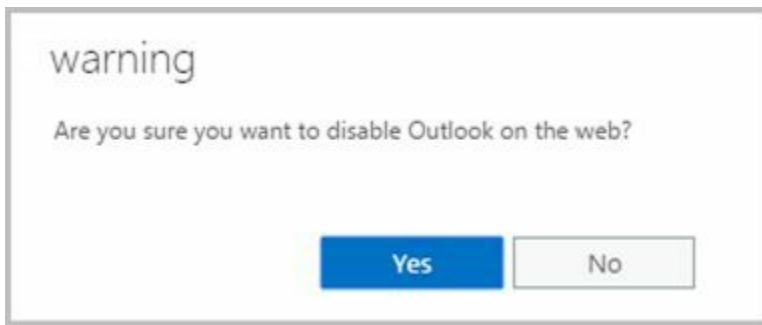
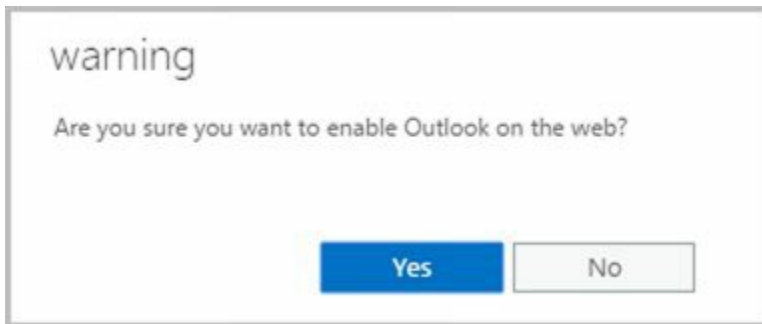


FIGURE 10-2 Use the options under Email Connectivity to manage a user's web access settings.

- To disable Outlook Web App for the user you selected, click **Disable**. When prompted to confirm, click **Yes**.




- To enable Outlook Web App for the user you selected, click Enable. When prompted to confirm, click Yes.



While you are working with Outlook Web App, you may want to determine the mailbox policy currently being applied. To view or change a user's Outlook Web App mailbox policy, do the following:

- Click **View Details**. In the Outlook Web App Mailbox Policy dialog box, the currently assigned policy is listed or the policy entry is blank, which means the default policy is currently applied.
- To assign a different policy, click **Browse**. Select a policy to view its enabled features. When you've selected the policy you want to use, click **OK**, and then click **Save**.



Configuring Mail Support for Outlook

You can install Outlook as a client on a user's computer. This section looks at the following topics:

- [Understanding address lists, offline address books, and autodiscover](#)
- [Configuring Outlook for the first time](#)
- [Adding Internet mail accounts to Outlook](#)
- [Reconfiguring Outlook mail support](#)

Unless specified otherwise, the procedures in this section work with desktop computers running current versions of Windows and Windows Server. Additionally, unless noted otherwise, the procedures work with Outlook 2010, Outlook 2013 and Outlook 2016.

Understanding Address Lists, Offline Address Books, and Autodiscover

Address lists are collections of recipients in an Exchange organization. Offline address books (OABs) are copies of address lists that are downloaded and cached on a computer so an Outlook user can access the address book while disconnected from the Exchange organization.

Every Exchange organization has a global address list and a default OAB. In the Exchange organization, address lists reside in Active Directory. If mobile users are disconnected from the Internet, they are unable to access the address lists stored on Exchange Online. If mobile users are disconnected from the corporate network, they are unable to access the address lists stored on Exchange 2016. To allow users to continue working when disconnected from the network, Exchange 2016 and Exchange Online generate offline address books and make them accessible to Outlook clients so that they can be downloaded and cached for use while working offline.

Although Exchange 2016 and Exchange Online continue to support public folders, public folders are not required for access to the global address list or the OAB. Exchange 2016 and Exchange Online provide these features through a web-based distribution point. Outlook clients use the web-based distribution point to obtain the global address list and the OAB automatically.

Exchange Online largely manages the default address lists and OABs automatically. On-premises Exchange, however, includes many configuration options, as discussed in the remainder of this section. For more information on global address lists and OABs, see "Managing Online Address Lists" and "Managing Offline Address Books" in Chapter 9 "Managing Addresses Online and Offline."

A designated Mailbox server, referred to as the *generation server*, is responsible for creating and updating the OABs. OAB data is produced by the Microsoft Exchange OABGen Service and stored in a special arbitration mailbox with the persisted

capability "OrganizationCapabilityOABGen." When a client initiates an OAB distribution request, the request is directed through a Mailbox server that routes the request to the Mailbox server hosting the OAB data. The OAB data is then distributed directly from the Mailbox server to the client.

Outlook 2010 and later as well as some mobile devices use the Autodiscover service to automatically configure themselves for access to Exchange. Outlook relies on DNS lookups to locate a host service (SRV) resource record for the Autodiscover service, then uses the user's credentials to authenticate to Active Directory and search for the Autodiscover connection points. After retrieving the connection points, the client connects to the first Mailbox server in the list and obtains the profile information. The connection point uses the globally unique identifier (GUID) for the user's mailbox plus the at symbol (@) and the domain portion of the user's primary SMTP address. The profile information includes the user's display name, the location of the user's mailbox server, connection settings for internal and external connectivity, MAPI over HTTP settings, and the URLs for Outlook features including those for free-busy data, the OAB, and Unified Messaging.

When you install a Mailbox server, an Autodiscover virtual directory is created on the default website in Internet Information Services (IIS), and an internal URL is set up for automatic discovery and other features, such as the OAB (which can be automatically discovered as well). Typically, the Autodiscover URL is either `https:// domain /autodiscover/autodiscover.xml` or `https://autodiscover. domain /autodiscover/autodiscover.xml`, where *domain* is your organization's primary SMTP domain address, such as `https://autodiscover.imagedlands.com/autodiscover/autodiscover.xml`. When you deploy multiple Mailbox servers, a connection point is created for each. This connection point stores the server's fully qualified domain name (FQDN) in the form `https:// servername /autodiscover/autodiscover.xml`, where *servername* is the FQDN of the Mailbox server, such as `https://server18.imagedlands.com/autodiscover/autodiscover.xml`.

The OAB virtual directory is the web-based distribution point for the OAB. By default, when you install a Mailbox server, this directory is created on the default website in IIS and configured for internal access. You can specify an external URL as well. Typically, the internal URL is set as `https:// servername /OAB`, where *servername* is the FQDN of the Mailbox server, such as `https://server18.imagedlands.com/OAB`.

For MAPI over HTTP to be automatically configured by using the Autodiscover service, external users running Outlook 2010 or later clients must have a valid Secure Sockets Layer (SSL) certificate on the Mailbox server that includes both the common name, such as `mail.imagedlands.com`, and a Subject Alternative name for the Autodiscover service, such as `autodiscover.imagedlands.com`. Also, the external URLs for the offline address book, Exchange Web Services, and MAPI over HTTP must be configured.

To configure the external URL for the OAB, you can use the `-ExternalUrl` parameter of

the Set-OABVirtualDirectory cmdlet. In the following example, you set the OAB external URL and configure it for use with SSL:

```
Set-OABVirtualDirectory -identity "Mailserver01\OAB (Default Web Site)"  
-externalurl https://mail.imagedlands.com/OAB -RequireSSL $true
```

To configure the external URL for Exchange Web Services, you can use the -ExternalUrl parameter of the Set-WebServicesVirtualDirectory cmdlet. The following example sets the Exchange Web Services external URL and configures it for use with basic authentication:

```
Set-WebServicesVirtualDirectory -identity "Mailserver01\EWS (Default Web Site)" -externalurl https://mail.imagedlands.com/EWS/Exchange.asmx  
-BasicAuthentication $True
```

To configure the external URL for MAPI over HTTP, you can use the -ExternalUrl parameter of the Set-MapiVirtualDirectory cmdlet. The following example sets the MAPI over HTTP external URL and configures it for use with NTLM and Negotiate authentication:

```
Set-MapiVirtualDirectory -identity "Mailserver01\mapi (Default Web Site)" -externalurl https://mail.imagedlands.com/mapi  
-IISAuthenticationMethods NTLM,Negotiate
```

If you want older clients to be able to use Outlook Anywhere, you can use the -ExternalHostname parameter of Set-OutlookAnywhere. The following example sets the external host name and configures authentication:

```
Set-OutlookAnywhere -Server Mailserver01 -ExternalHostname  
"mail.imagedlands.com" -ExternalClientAuthenticationMethod Negotiate  
-InternetClientAuthenticationMethod NTLM  
-IISAuthenticationMethods Basic, NTLM, Negotiate  
-SSLOffloading $False
```

Once you've configured these options, you can test the Availability service, Outlook Anywhere, and the Offline Address Book service by using Test-OutlookWebServices. Here are examples:

```
Test-OutlookWebServices -ClientAccessServer "Mailserver01"  
Test-OutlookWebServices -Identity "willams@imagedlands.com"
```

Use Test-MapiConnectivity to test MAPI over HTTP. Here are examples:

```
Test-MapiConnectivity -Server "Mailserver01"  
Test-MapiConnectivity -Identity "willams@imagedlands.com"
```

Configuring Outlook for the First Time

You can install Outlook as a standalone product or as part of Microsoft Office. Outlook can be used to connect to the following types of email servers:

- **Microsoft Exchange** Connects directly to Exchange Server, Exchange Online, or

both; best for users who are connected to the organization's network. Users will have full access to Exchange. If users plan to connect to Exchange using MAPI over HTTP or Outlook Anywhere (RPC over HTTP), this is the option to choose as well. With Exchange, users can check mail on an email server and access any private or public folders to which they have been granted permissions. If you define a personal folder and specify that new email messages should be delivered to it, messages can be delivered to a personal folder on a user's computer.

- **POP3** Connects to Exchange 2016 or another POP3 email server through the Internet; best for users who are connecting from a remote location, such as a home or a remote office, using dial-up or broadband Internet access. With POP3, users can check mail on an email server and download it to their inboxes. Users can't, however, synchronize mailbox folders or access private or public folders on the server. By using advanced configuration settings, the user can elect to download the mail and leave it on the server for future use. By leaving the mail on the server, the user can check mail in Outlook Web App or on a home computer and then still download it to an office computer later.
- **IMAP4** Connects to Exchange 2016 or another IMAP4 email server through the Internet; best for users who are connecting from a remote location, such as a home or a remote office, using dial-up or broadband Internet access. Also well suited for users who have a single computer, such as a laptop, that they use to check mail both at the office and away from it. With IMAP4, users can check mail on an email server and synchronize mailbox folders. Users can also download only message headers and then access each message individually to download it. Unlike POP3, IMAP4 has no option to leave mail on the server. IMAP4 also lets users access public and private folders on an Exchange server.
- **ActiveSync** Connects to an Exchange ActiveSync compatible service, such as Outlook.com, through the Internet; best as an additional email configuration option. Users can have an external email account with a web-based email service that they can check in addition to corporate email.
- **Additional server types** Connects to a third-party mail server or other services, such as Outlook Mobile Text Messaging. If your organization has multiple types of mail servers, including Exchange Server, you'll probably want to configure a connection to Exchange Server first and then add more email account configurations later.

To begin, log on to the computer as the user whose email you are configuring or have the user log on. If the computer is part of a domain, log on using the user's domain account. If you are configuring email for use with a direct Exchange 2016 or Exchange Online connection rather than a POP3, IMAP4, or ActiveSync connection, ensure that the user's mailbox has been created. If the user's mailbox has not been created, auto-setup will fail, as will the rest of the account configuration.

The first time you start Outlook, the application runs the Welcome Wizard. You can use the Welcome Wizard to configure email for Exchange, POP3, IMAP4, and ActiveSync mail servers, as discussed in the sections that follow.

First-Time Configuration: Connecting to Exchange Server

With Outlook 2010 or later, you can use the Welcome Wizard to configure email for Exchange 2016 or Exchange Online in Outlook by completing the following steps:

1. Start Outlook and click **Next** on the Welcome page. The procedure is nearly identical whether you are working with Outlook 2010, Outlook 2013 or Outlook 2016.
2. When prompted to indicate whether you would like to configure an email account, verify that **Yes** is selected and then click **Next**.



Use Outlook to connect to email accounts, such as your organization's Microsoft Exchange Server or an Exchange Online account as part of Microsoft Office 365. Outlook also works with POP, IMAP, and Exchange ActiveSync accounts.

Do you want to set up Outlook to connect to an email account?

Yes
 No

3. The next page of the wizard varies depending on the computer's current configuration:
- For computers that are part of a domain and for users that have an existing Exchange Server mailbox, the wizard uses the Autodiscover feature to automatically discover the required account information.



E-mail Account

Your Name:
Example: Ellen Adams

E-mail Address:
Example: ellen@contoso.com

- For computers that are part of a domain and for users without an on-premises Exchange mailbox, leave the wizard open, create the user's Exchange mailbox, and then proceed with the wizard once the mailbox is automatically discovered.
- For all other configurations, including computers that are part of a workgroup, and computers on which you are logged on locally, Outlook assumes you want to configure an Internet email account for the user. Enter the user's account name and email address. Then type and confirm the user's password (see Figure 10-3).

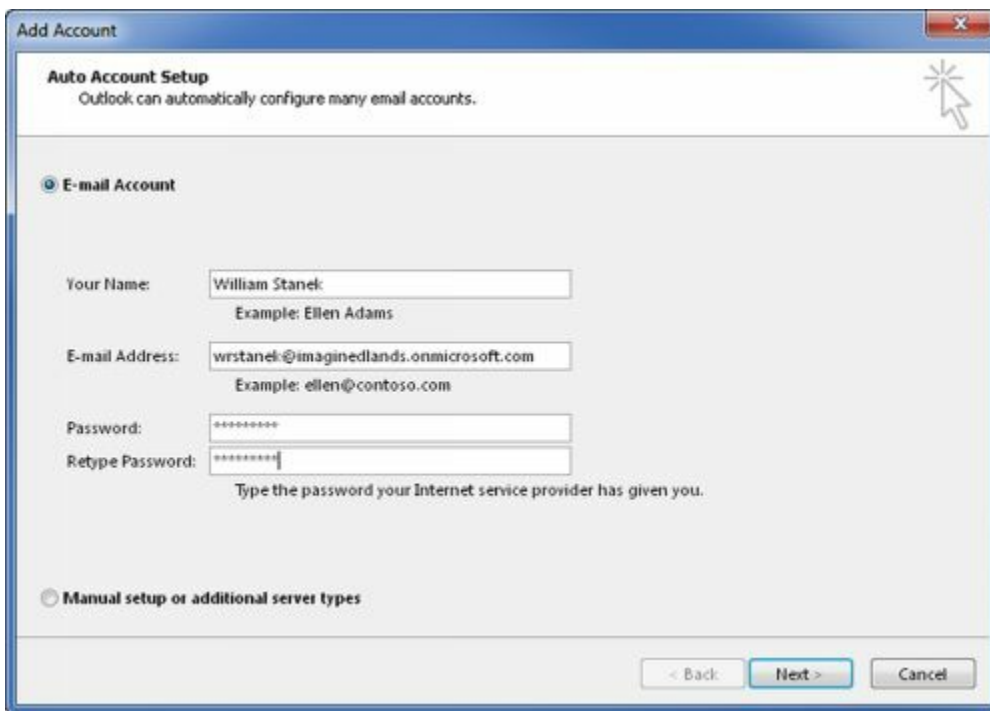
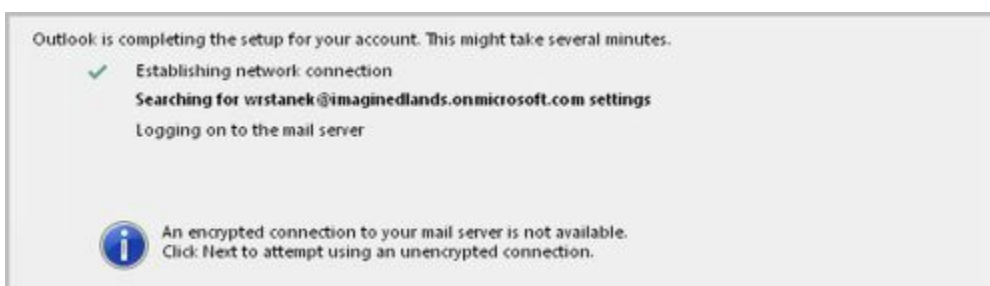
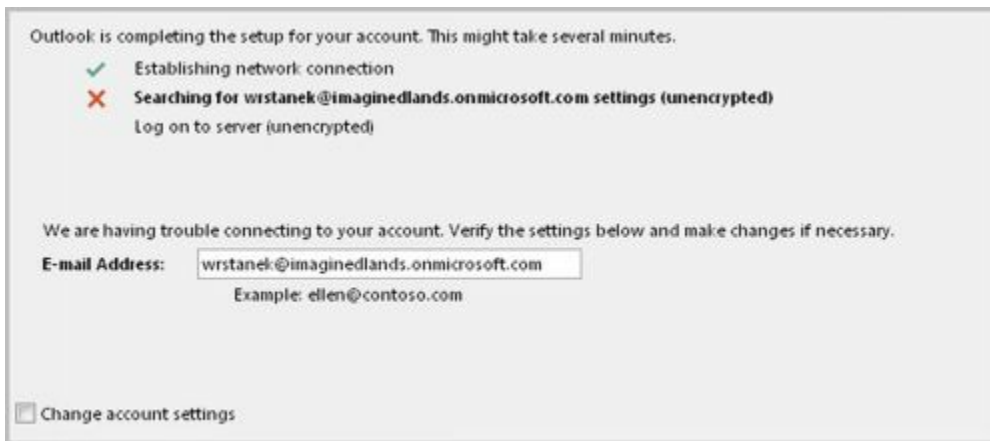


FIGURE 10-3 Although the Wizard can automatically fill in account information when you are logged on to a domain, the wizard does not do this for other configurations.

4. After you click **Next**, the wizard uses the new **Auto Account Setup** feature to automatically discover the rest of the information needed to configure the account and then uses the settings to log on to the server. If the auto-configuration and server logon are successful, click **Finish** and skip the remaining steps in this procedure. The wizard then sets up the user's Exchange mailbox on the computer as appropriate.
5. If auto-configuration is not successful, click **Next** so that the wizard can attempt to establish an unencrypted connection to the server. If the auto-configuration and server logon are successful this time, click **Finish** and then skip the remaining steps in this procedure.



6. If auto-configuration fails twice, you'll see a prompt to confirm the user's email address. If the email address is incorrect, correct it, and then click **Retry**. If the auto-configuration and server logon are successful this time, click **Finish** and then skip the remaining steps in this procedure.



7. If all attempts at auto-configuration fail, you can try to configure settings manually (and might also want to confirm that the Autodiscover service is working properly). Click **Next** . On the Choose Service page, select a service. Click **Next** . On the next wizard page, complete the necessary information for the type of email service you selected. If necessary, click **More Settings** . Use the Properties dialog box to configure the additional required settings and then click **OK** . Click **Next** and then click **Finish** to complete the mail configuration.

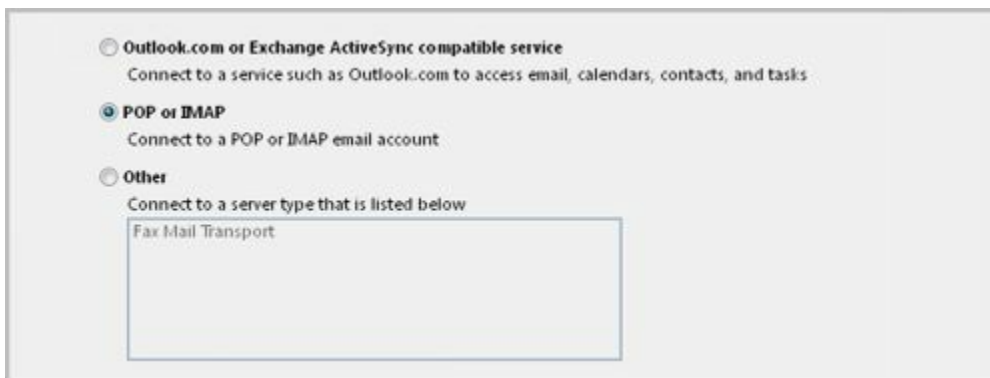
First-Time Configuration: Connecting to Internet Email Servers

When a user is logged on to a domain, Outlook automatically attempts to configure itself for use with the user's Exchange mailbox as part of its initial configuration. This configuration works for internal users but not for remote users who need or prefer to access Exchange using POP3 or IMAP4 (rather than MAPI over HTTP or Outlook Anywhere [which uses RPC over HTTP]). For these users, you can complete the first-time configuration of Outlook by following these steps:

1. In the Welcome Wizard, prompted to indicate whether you would like to configure an email account, verify that **Yes** is selected and then click **Next** .
2. Select the manual setup option. In Outlook 2010, this checkbox is labeled as **Manually Configure Server Settings Or Additional Server Types**. In Outlook 2013 and Outlook 2016, this checkbox is labeled as **Manual Setup Or Additional Server Types**. Click **Next** .

A screenshot of the Outlook manual setup wizard. At the top, there is a radio button labeled "E-mail Account" which is selected. Below this, there are four input fields: "Your Name:" with a text box and an example "Example: Ellen Adams"; "E-mail Address:" with a text box and an example "Example: ellen@contoso.com"; "Password:" with a text box; and "Re-type Password:" with a text box. Below the password fields, there is a note: "Type the password your Internet service provider has given you." At the bottom, there is a radio button labeled "Manual setup or additional server types" which is selected.

3. On the Choose Service page, choose the service to use. In Outlook 2010, choose Internet E-Mail as the service. In Outlook 2013 and Outlook 2016, choose POP Or IMAP as the service. Click **Next** .



Outlook.com or Exchange ActiveSync compatible service
Connect to a service such as Outlook.com to access email, calendars, contacts, and tasks

POP or IMAP
Connect to a POP or IMAP email account

Other
Connect to a server type that is listed below

Fax Mail Transport

4. In the Your Name text box, type the name to appear in the From field of outgoing messages for this user, such as **William Stanek** .
5. In the E-Mail Address text box, type the email address of the user. Be sure to type the email user name as well as the domain name, such as **williams@imaginedlands.com** .
6. From the Account Type list, select POP3 or IMAP4 as the type of protocol to use for the incoming mail server. The advantages and disadvantages of these protocols are as follows:
 - POP3 is used to check mail on an email server and download it to the user's inbox. The user can't access private or public folders on the server. By using advanced configuration settings, the user can elect to download email and leave it on the server for future use. By leaving the email on the server, the user can check a message on a home computer and still download it to an office computer later.
 - IMAP4 is used to check mail on an email server and download message headers. The user can then access each email individually and download it. Unlike POP3, IMAP4 has no option to leave mail on the server. IMAP4 also lets users access public and private folders on an Exchange server. It is best suited for users who have a single computer, such as a laptop, that they use to check mail both at the office and away from it.



User Information

Your Name: William Stanek

Email Address: williams@imaginedlands.co

Server Information

Account Type: POP3

Incoming mail server: pop3.imaginedlands.com

Outgoing mail server (SMTP): smtp.imaginedlands.com

Logon Information

User Name: williams@imaginedlands.co

Password: *****

Remember password

Require logon using Secure Password Authentication (SPA)

Test Account Settings

We recommend that you test your account to ensure that the entries are correct.

Test Account Settings ...

Automatically test account settings when Next is clicked

Deliver new messages to:

New Outlook Data File

Existing Outlook Data File

Browse

More Settings ...

7. Enter the FQDN for the incoming and outgoing mail servers. Although these

entries are often the same, some organizations have different incoming and outgoing mail servers. If you are not certain of your mail servers' FQDN, contact your network administrator.

NOTE If you're connecting to Exchange with POP3 or IMAP4, you should enter the FQDN for the Exchange server rather than just the host name. For example, you would use MailServer.imaginedlands.com instead of MailServer. This ensures Outlook will be able to find the Exchange server.

8. Under Logon Information, type the user's logon name and password. If the mail server requires secure logon, select the Require Logon Using Security Password Authentication check box.
9. To verify the settings, click **Test Account Settings** . Outlook verifies connectivity to the Internet and then logs on to the Mail server. Next, Outlook sends a test message to the specified mail server. If the test fails, note the errors and make corrections as necessary.
10. If necessary, click More Settings . Use the Properties dialog box to configure the additional required settings and then click **OK** . When you are ready to continue, click **Next** , and then click **Finish** to complete the configuration.

Configuring Outlook for Exchange

If you didn't configure Outlook to use Exchange the first time it was started and elected to use Outlook without an email account, don't worry: You can change the Outlook configuration to use Exchange. It does take a bit of extra work, however.

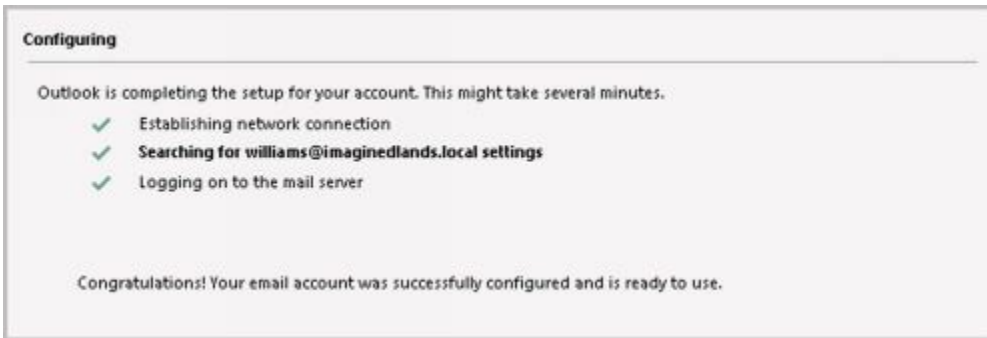
Follow these steps to configure Outlook to use Exchange:

1. In Outlook 2013 or Outlook 2016, click **File** and then select **Add Account**. If you are using Outlook 2010, you must select Tools, Account Settings, New and then select Microsoft Exchange as the e-mail service.
2. Outlook assumes you want to configure an Internet email account for the user. You will, however, provide the information needed for Exchange. Enter the user's account name and email address for Exchange. Then type and confirm the user's domain password.



3. Click **Next**. The wizard uses the new Auto Account Setup feature to automatically discover the rest of the information needed to configure the account

and then uses the settings to log on to the server. If the auto-configuration and server logon are successful, click **Finish**.



Adding Internet Mail Accounts to Outlook

Through email account configuration, each mail profile for Outlook supports only one Exchange Server account at a time. If you need access to multiple Exchange Server mailboxes in the same mail profile, you must configure access to these mailboxes as discussed in the section “Accessing Multiple Exchange Mailboxes” later in the chapter.

Although you can configure only one Exchange email account for each mail profile, Outlook allows you to retrieve mail from both Exchange Online and Exchange Server as well as from multiple Internet servers. For example, you can configure Outlook to check mail on the corporate Exchange server, a personal account with an ActiveSync compatible service, and Exchange Online.

You can add Internet mail accounts to Outlook. In Outlook, complete the following steps:

1. In Outlook 2013 or Outlook 2016, click **File** and then select **Add Account**. If you are using Outlook 2010, you must select **Tools, Account Settings, New** and then select **POP3, IMAP, Or HTTP** as the e-mail service.
2. Click **Next**. The wizard tries to use the new Auto Account Setup feature to automatically discover the rest of the information needed to configure the account and then uses the settings to log on to the server.
3. If the auto-configuration and server logon are successful, click **Finish**. **Otherwise**, follow steps 2–10 outlined previously in the “First-time Configuration: Connecting to Internet Email Servers” section.

Repairing and Changing Outlook Mail Accounts

When you first configure Outlook on a computer, you can configure it to connect to an Exchange server, to Exchange Online, to Internet email, or to another email server. With Exchange Server, Outlook can use MAPI over HTTP or RPC over HTTP to connect to the appropriate Mailbox server and access the appropriate mailbox. If a user’s mailbox is moved to a different server within the Exchange organization, the user is connected to this server automatically the next time he or she starts Outlook. If, for some reason, a user has a problem connecting to Exchange or needs to update configuration settings, you can use a repair operation. Repairing the user’s account restarts the Auto Account Setup feature.

With non-Exchange servers, access to email very much depends on the account and server configuration remaining the same. If the account or server configuration changes, the account configuration in Outlook must be updated. The easiest way to do this is with a repair operation.

To start a repair, follow these steps:

1. Log on as the domain account of the user for whom you are repairing email.
2. In Outlook 2010, click the Office button, click **Account Settings** , and then select the **Account Settings** option. In Outlook 2013 or Outlook 2016, on the File pane, click **Account Settings** , and then select the **Account Settings** option.
3. In the Account Settings dialog box, the E-Mail tab lists all currently configured email accounts by name. Select the account to repair and then click **Repair** .
4. On the Auto Account Setup page, check the account settings. With Exchange accounts for domain users and with Exchange Online, you cannot change the displayed information. With other accounts, you can modify the user's email address and password, as necessary.
5. When you click **Next** , the Repair E-Mail Account Wizard contacts the mail server and tries to determine the correct account settings. If the auto-configuration and server logon are successful, click **Finish** . Skip the remaining steps in this procedure.
6. If auto-configuration is not successful, click **Next** so that the wizard can attempt to establish an unencrypted connection to the server. If the auto-configuration and server logon are successful this time, click **Finish** and then skip the remaining steps in this procedure. You must restart Outlook.

NOTE You may be prompted to confirm the user's credentials. If so, type the user's password, select the Remember My Credentials checkbox, and then click **OK**.

7. If auto-configuration fails twice, you can try to configure settings manually. Select the manual setup option, and then click **Next**.
8. Use the fields provided to update the mail account configuration. If you need to configure additional settings beyond the user, server, and logon information, click the More button (**...**), and then use the Properties dialog box to configure the additional required settings. When you are finished, click **OK** to close the Properties dialog box.
9. To check the new settings, click **Test Account Settings** .
10. Click **Next** , and then click **Finish** .

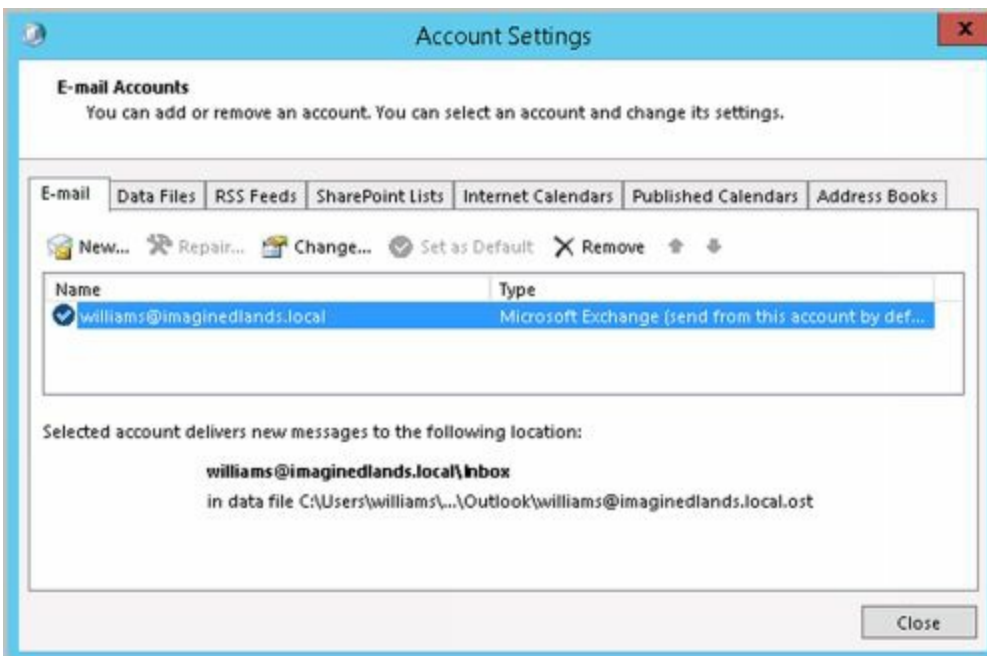
In some cases, if you've incorrectly configured Exchange, you might not be able to start Outlook and access the Account Settings dialog box. In this case, you can repair the settings using the following procedure:

1. Start the Mail utility. In Control Panel, click **Small Icons** on the View By list and

then start the Mail app by clicking its icon or by double-clicking its icon.



2. In the Mail Setup–Outlook dialog box, click **E-Mail Accounts** . The Accounts Settings dialog box appears.



3. In the Account Settings dialog box, the E-Mail tab is selected by default. Click the incorrectly configured Exchange account and then do one of the following:
 - Click **Change** to modify the Exchange settings using the techniques discussed previously.
 - Click **Remove** to remove the Exchange settings so that they are no longer used by Outlook.
4. When you are finished, close the Mail Setup–Outlook dialog box, and then start Outlook.

For POP3 or IMAP4, you can change a user's email configuration at any time by completing the following steps:

1. In Outlook 2010, click the Office button, click **Account Settings** , and then select the **Account Settings** option. In Outlook 2013 or Outlook 2016, on the File pane, click **Account Settings** , and then select the **Account Settings** option.

2. In the Account Settings dialog box, the E-Mail tab lists all currently configured email accounts by name. Select the account you want to work with, and then click **Change** .
3. Use the fields provided to update the mail account configuration. If you need to configure additional settings beyond the user, server, and logon information, click the More button (**...**), and then use the Properties dialog box to configure the additional required settings. When you are finished, click **OK** to close the Properties dialog box.
4. To check the new settings click **Test Account Settings** .
5. Click **Next** , and then click **Finish** .

Leaving Mail on the Server with POP3

If the user connects to an Internet e-mail server, an advantage of POP3 is that it lets a user leave mail on the server. By doing this, the user can check mail on a home computer and still download it to an office computer later.

With Outlook, you can configure POP3 accounts to leave mail on the server by completing the following steps:

1. In Outlook 2010, click the Office button, click **Account Settings** , and then select the **Account Settings** option. In Outlook 2013 or Outlook 2016, on the File pane, click **Account Settings** , and then select the **Account Settings** option.
2. In the Account Settings dialog box, select the POP3 mail account you want to modify and then click **Change** .
3. Click the More button (**...**)to display the Internet E-Mail Settings dialog box.
4. In the Internet E-Mail Settings dialog box, click the **Advanced** tab, as shown in Figure 10-5.
5. Use the options below Delivery to configure how and when mail should be left on the server. To enable this option, select the **Leave A Copy Of Messages On The Server** check box. The additional options depend on the client configuration. Options you might see include the following:
 - **Remove From Server After NDays** Select this option if the user will be connecting to an Internet service provider (ISP) and you want to delete messages from the server after a specified number of days. By deleting ISP mail periodically, you ensure that the mailbox size doesn't exceed the limit.

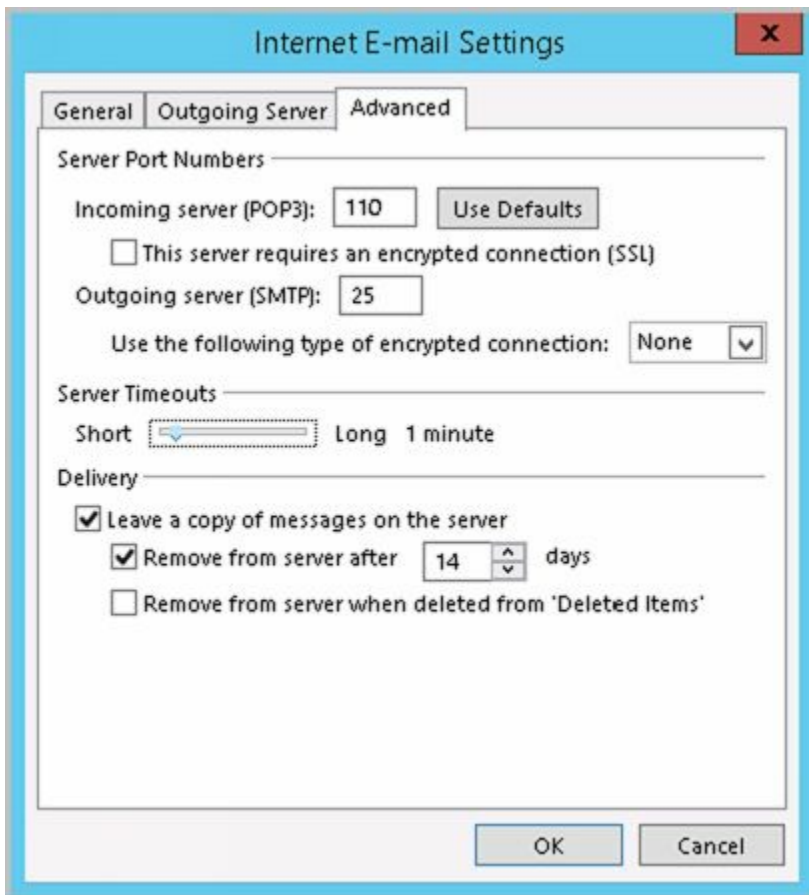


FIGURE 10-5 Using the Advanced tab to configure how and when mail should be left on the server.

- **Remove From Server When Deleted From “Deleted Items”** Select this option to delete messages from the server when the user deletes them from the Deleted Items folder. You’ll see this option with Internet-only Outlook configurations.
6. Click **OK** when you’ve finished changing the account settings.
 7. Click **Next** , and then click **Finish** . Click **Close** to close the Account Settings dialog box.

Checking Private and Public Folders with IMAP4 and UNIX Mail Servers

With IMAP4, you can check public and private folders on a mail server. This option is enabled by default, but the default settings might not work properly with UNIX mail servers.

With Outlook, you can check or change the folder settings used by IMAP4 by completing the following steps:

1. Start Outlook. In Outlook 2010, click the Office button, click **Account Settings** , and then select the **Account Settings** option. In Outlook 2013, on the File pane, click **Account Settings** , and then select the **Account Settings** option.
2. In the Account Settings dialog box, select the IMAP4 mail account you want to modify and then click **Change** .
3. Click the More button (**...**) to display the Internet E-Mail Settings dialog box.
4. In the Internet E-Mail Settings dialog box, click the **Advanced** tab, as shown in Figure 10-6.
5. If the account connects to a UNIX mail server, enter the path to the mailbox folder on the server, such as **~williams/mail** —don't end the folder path with a forward slash (/)—and then click **OK** .
6. Click **Next** , and then click **Finish** .

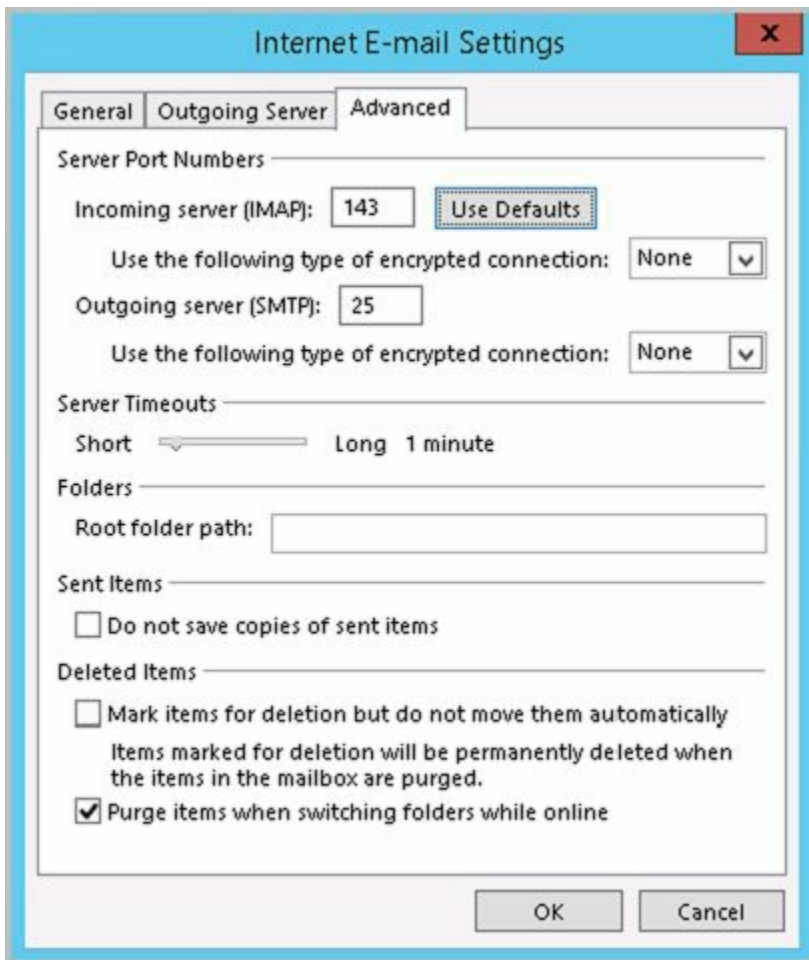


FIGURE 10-6 Using the Advanced tab to configure how folders are used with IMAP4 mail accounts.

Managing the Exchange Configuration in Outlook

Whenever you use Outlook to connect to Exchange, you have several options for optimizing the way mail is handled. These options include the following:

- [Email delivery and processing](#)
- [Remote mail](#)
- [Scheduled connections](#)
- [Multiple mailboxes](#)

Each of these options is examined in this section.

Managing Delivery and Processing Email Messages

When Outlook uses Exchange, you have strict control over how email is delivered and processed. Exchange mail can be delivered in one of two ways:

- [To server mailboxes with local copies](#)
- [To personal folders](#)

Exchange mail can be processed by any of the information services configured for use in Outlook. These information services include the following:

- [Microsoft Exchange](#)
- [Internet email](#)

Let's look at how you use each of these delivery and processing options.

Using Server Mailboxes

When you are using Outlook 2010 or later with Exchange 2016 or Exchange Online, server mailboxes with local copies are the default configuration option. With server mailboxes, new email is delivered to a mailbox on the Exchange server, and users can view or receive new mail only when they're connected to Exchange. When users are connected to Exchange, Outlook retrieves their mail and stores a local copy on their computer in addition to the email stored on Exchange.

The local copy of a user's mail is stored in an offline folder .ost file. With Windows 7 and later, the default location of a .ost file is `%LocalAppData%\Microsoft\Outlook`, where `%LocalAppData%` is a user-specific environment variable that points to a user's local application data. Using server mailboxes offers users protected storage and the ability to have a single point of recovery in case something happens to their computer.

Using Personal Folders

An alternative to using server mailboxes is to use personal folders. Personal folders are stored in a .pst file on the user's computer. With personal folders, you can specify that mail should be delivered to the user's inbox and stored on the server or that mail should be delivered only to the user's inbox. Users have personal folders when Outlook is

configured to use Internet email or other email servers. Users might also have personal folders if the auto-archive feature is used to archive messages.

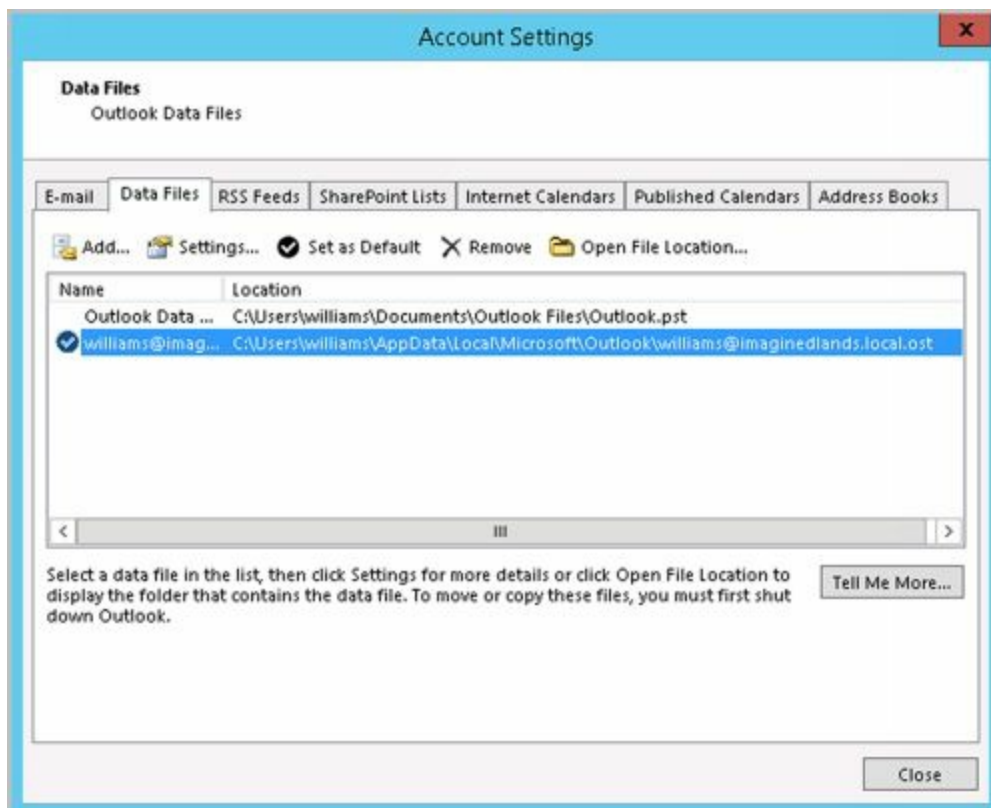
REAL WORLD With Windows 7 and later, the default location of a .pst file is `%LocalAppData%\Microsoft\Outlook`, where `%LocalAppData%` is a user-specific environment variable that points to a user's local application data. Personal folders are best suited for mobile users who check mail through dial-up connections and who might not be able to use a dial-up connection to connect directly to Exchange.

Users with personal folders lose the advantages that server-based folders offer—namely, protected storage and the ability to have a single point of recovery in case of failure. In addition, .pst files have many disadvantages. They get corrupted more frequently and, on these occasions, you must use the Inbox Repair Tool to restore the file. If the hard disk on a user's computer fails, you can recover the mail only if the .pst file has been backed up. Unfortunately, most workstations aren't backed up regularly (if at all), and the onus of backing up the .pst file falls on the user, who might or might not understand how to do this.

Determining the Presence of Personal Folders

You can determine the presence of personal folders by following these steps:

1. Start Outlook. In Outlook 2010, click the Office button, click **Account Settings**, and then select the **Account Settings** option. In Outlook 2013 or Outlook 2016, on the File pane, click **Account Settings**, and then select the **Account Settings** option.



2. In the Account Settings dialog box, click the **Data Files** tab.

3. The location of the data file associated with each email account is listed. If the file name ends in .pst, the account is using a personal folder.

Creating New or Opening Existing Personal Folders

If personal folders aren't available and you want to configure them, follow these steps:

1. Start Outlook. In Outlook 2010, click the Office button, click **Account Settings**, and then select the **Account Settings** option. In Outlook 2013 or Outlook 2016, on the File pane, click **Account Settings**, and then select the **Account Settings** option.
2. In the Account Settings dialog box, click the **Data Files** tab.
3. Click **Add**. If the New Outlook Data File dialog box appears, Office Outlook Personal Folders File (.pst) should be selected by default. Click **OK**.
4. Use the Create Or Open Outlook Data File dialog box, as shown in Figure 10-7, to create a new .pst file or open an existing .pst file:

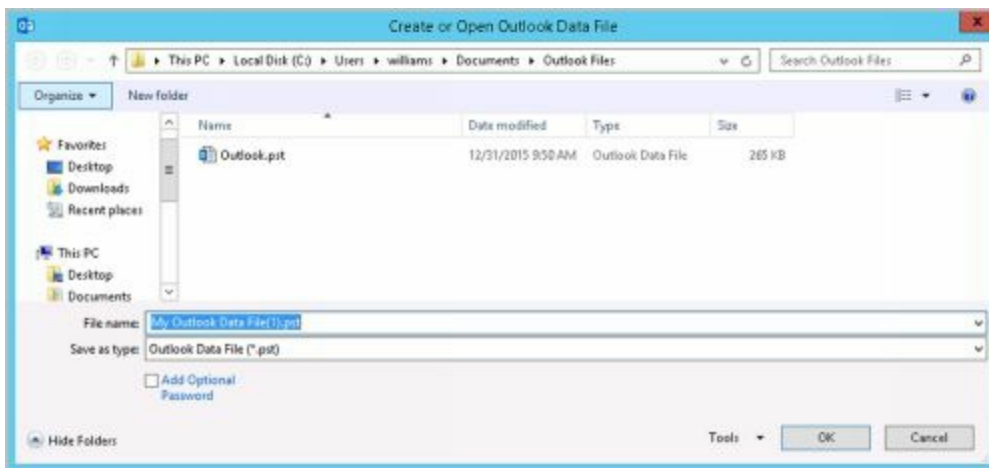


FIGURE 10-7 Using the Create Or Open Outlook Data File dialog box to search for an existing .pst file or to create a new one.

- To create a new .pst file in the default folder, type a name for the Outlook data file in the text box provided or accept the default value. To secure the file and ensure only a person with this password can access the file, select the Add Optional Password checkbox. In the Create Microsoft Personal Folders dialog box, specify a password, verify a password for the .pst file, and click OK.
- To create a new .pst file in a nondefault folder, click Browse Folders to show the folder view if it is hidden. Browse for the folder you want to use, type the file name in the text box provided or accept the default value, and then click OK. Optionally, select the Add Optional Password checkbox. In the Create Microsoft Personal Folders dialog box, specify a password, verify a password for the .pst file, and click OK.
- To open an existing .pst file, click Browse Folders to show the folder view if it is hidden. Browse to the folder containing the .pst file. Select the .pst file, and then click OK. In the Personal Folders dialog box, use the options provided to change the current password or compact the personal folder, and then click OK.

NOTE It is important to be aware that Exchange Server does not ship with any password recovery utility for .pst files. If a user sets a password on a .pst file and forgets it, the Exchange administrator has no way to reset it. You might find third-party vendors who make password-cracking or recovery tools, but they are not guaranteed to work and they are not supported by Microsoft.

5. Click **Close** . The personal folder you've selected or created is displayed in the Outlook folder list. You should see related subfolders as well.

Delivering Mail to Personal Folders

When you configure mail to be delivered to a personal folder, Outlook saves email messages only locally on the computer. As a result, Outlook removes the messages from Exchange Server after delivery and you can access the messages only on the currently logged-on computer.

If you want mail to be delivered to a personal folder, complete the following steps:

1. Start Outlook. In Outlook 2010, click the Office button, click **Account Settings** , and then select the **Account Settings** option. In Outlook 2013 or Outlook 2016, on the File pane, click **Account Settings** , and then select the **Account Settings** option.
2. In the Account Settings dialog box, click the **Data Files** tab.
3. Select the .pst file to use in the list of data files provided, and then click **Set As Default** .
4. When prompted to confirm, click **Yes** and then click **Close** .
5. Exit and restart Outlook. Outlook will now use personal folders.

If you want mail to resume using server-stored mail, complete the following steps:

1. Start Outlook. In Outlook 2010, click the Office button, click **Account Settings** , and then select the **Account Settings** option. In Outlook 2013 or Outlook 2016, on the File pane, click **Account Settings** , and then select the **Account Settings** option.
2. In the Account Settings dialog box, click the **Data Files** tab.
3. Select the .ost file to use in the list of data files provided, and then click **Set As Default** .
4. When prompted to confirm, click **OK** and then click **Close** .
5. Exit and restart Outlook. Outlook will now use personal folders.

Repairing .pst data files

When Outlook uses personal folders, you can use the Inbox Repair tool (scanpst.exe) to analyze and repair corrupted data files.

- With Office 2010 and Office 2013, this tool is stored in the %SystemDrive%\Program Files\Microsoft Office\Office *Version* folder, where *Version* is the internal version of

Office you are using, such as Office15 for Outlook 2013.

- With Office 2016, this tool is stored in the %SystemDrive%\Program Files\Microsoft Office\root\Office16 folder, such as c:\program files (x86)\Microsoft Office\root\Office16.

If a .pst file won't open or is damaged, you can use the Inbox Repair tool to repair it by completing the following these steps:

1. Exit Outlook. Open the Office folder in File Explorer and then double-click the Inbox Repair tool (scanpst.exe).



2. Click **Browse** . In the Select File To Scan dialog box, browse to the folder where .pst files are stored, select the .pst file you want to work with, and then click **Open** . Generally, .pst files are either stored in %LocalAppData%\Microsoft\Outlook, where %LocalAppData% is a user-specific environment variable that points to a user's local application data, or in the %UserProfile%\Documents\Outlook Files folder, where %UserProfile% is a user-specific environment variable that points to the user's local profile data.
3. Click **Start** , and the Inbox Repair tool will begin analyzing the file. The larger the file the longer the analysis will take.
4. If errors are found, click **Repair** to start the repair process. The Inbox Repair tool will create a copy of the .pst file before attempting the repair operation. During the repair, the Inbox Repair tool will rebuild the .pst file. This backup will be stored in the same folder as the original .pst file.



5. Start Outlook with the profile that contains the .pst file that you repaired. Press Ctrl+6 to display the Folder List view and look for a folder named Recovered Personal Folders. This folder contains the default Outlook folders as well as a Lost And Found folder, which contains any items recovered by the Inbox Repair tool.

NOTE You can also display the Folders List view by clicking the More button (**•••**) in the Navigation menu and then selecting Folders.

6. Create a new .pst data file to store your mail items. Drag the items from the Lost And Found folder into the appropriate folder under the new Personal folders. When you've moved all the items, you can remove the Recovered Personal Folders .
7. The Inbox Repair tool creates a backup of the original .pst file and names it with the .bak file extension. By default this file is stored in the same location as the original .pst file. If you make a copy of this file and name it with a .pst extension, you may be able to recover additional items. To do this, add the .pst file to the mail profile and then move any additional mail items from this old .pst file to the new data file created in step 6.

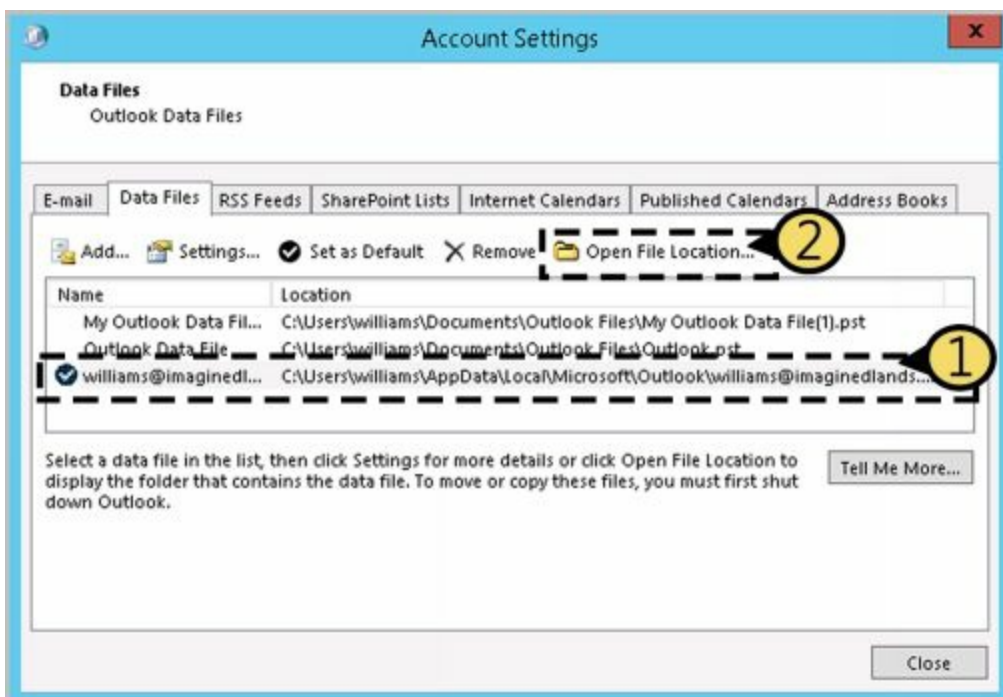
Repairing .ost data files

When Outlook uses server mailboxes, .ost data files contain copies of information saved on the server. If an .ost file won't open or is damaged, you can re-create the file by completing the following these steps:

1. Exit Outlook. Start the Mail utility. Press the Windows key +I and then click **Control Panel** . In Control Panel, click **Small Icons** on the View By list and then start the Mail app by double-clicking its icon.



2. In the Mail Setup–Outlook dialog box, click **Data Files** . This opens the Account Settings dialog box with the Data Files tab selected.



3. Select the Exchange account and then click **Open File Location** . This opens File Explorer to the location of the data file. Note this location. By default, .ost files are stored in *%LocalAppData%\Microsoft\Outlook*, where *%LocalAppData%* is a user-specific environment variable that points to a user’s local application data.
4. Close the Account Settings and Mail Setup dialog boxes. In File Explorer, right-click the .ost file and then click **Delete** . If you are unable to delete the file, make sure all mail and Office windows are closed.
5. Start Outlook. Download a copy of the mail items again to automatically re-create the .ost file.

Accessing Multiple Exchange Mailboxes

Earlier in the chapter, I discussed how users could check multiple Internet mail accounts in Outlook. You might have wondered whether users could check multiple Exchange mailboxes as well—and they can. Users often need to access multiple Exchange

mailboxes for many reasons:

- Help desk administrators might need access to the help desk mailbox in addition to their own mailboxes.
- Managers might need temporary access to the mailboxes of subordinates who are on vacation.
- Project team members may need to access mailboxes set up for long-term projects.
- Resource mailboxes might need to be set up for accounts payable, human resources, corporate information, and so on.

Normally, a one-to-one relationship exists between user accounts and Exchange mailboxes. You create a user account and add a mailbox to it; only this user can access the mailbox directly through Exchange. To change this setup, you must change the permissions on the mailbox. One way to change mailbox access permissions is to do the following:

1. Log on to Exchange as the owner of the mailbox.
2. Delegate access to the mailbox to one or more additional users.
3. Have users with delegated access log on to Exchange and open the mailbox.

The sections that follow examine each of these steps in detail.

Logging on to Exchange as the Mailbox Owner

Logging on to Exchange as the mailbox owner allows you to delegate access to the mailbox. Before you can do this, however, you must complete the following steps:

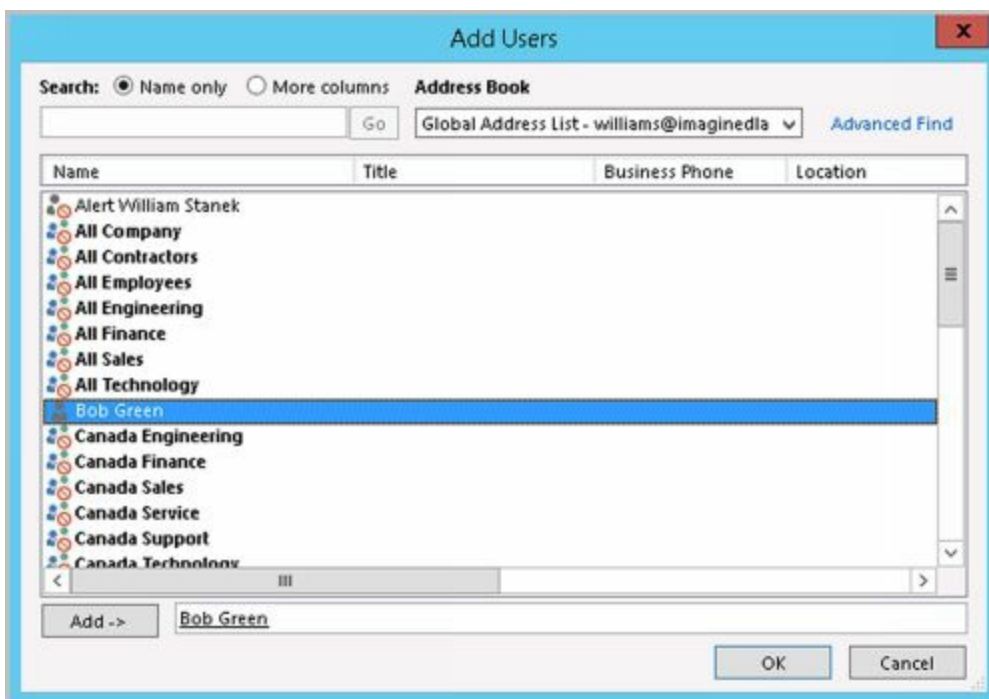
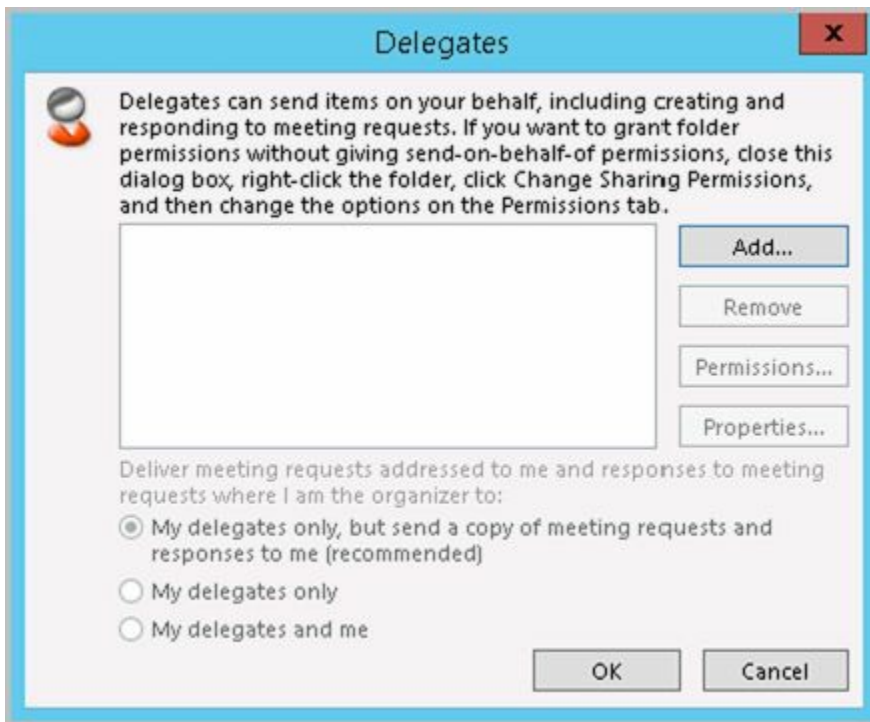
1. Log on as the user or have the user log on for you.
2. Start Outlook. Make sure that mail support is configured to use server mailboxes. If necessary, configure this support, which creates the mail profile for the user.
3. After you configure Outlook to use Exchange, you should be able to log on to Exchange as the mailbox owner.

TIP With multiple mailbox users, you should configure the mailbox to deliver mail to the server rather than to a personal folder. In this way, the mail can be checked by one or more mailbox users.

Delegating Mailbox Access

After you've logged on as the mailbox owner, you can delegate access to the mailbox by completing these steps:

1. Start Outlook. Open the Delegates dialog box by doing one of the following:
 - In Outlook 2010, click the Office button, click **Account Settings**, and then select the **Account Settings** option. On the Delegates tab or in the Delegates dialog box, click **Add**.
 - In Outlook 2013 or Outlook 2016, on the File pane, click **Account Settings**, and then select the **Delegate Access** option. In the Delegates dialog box, click **Add**.



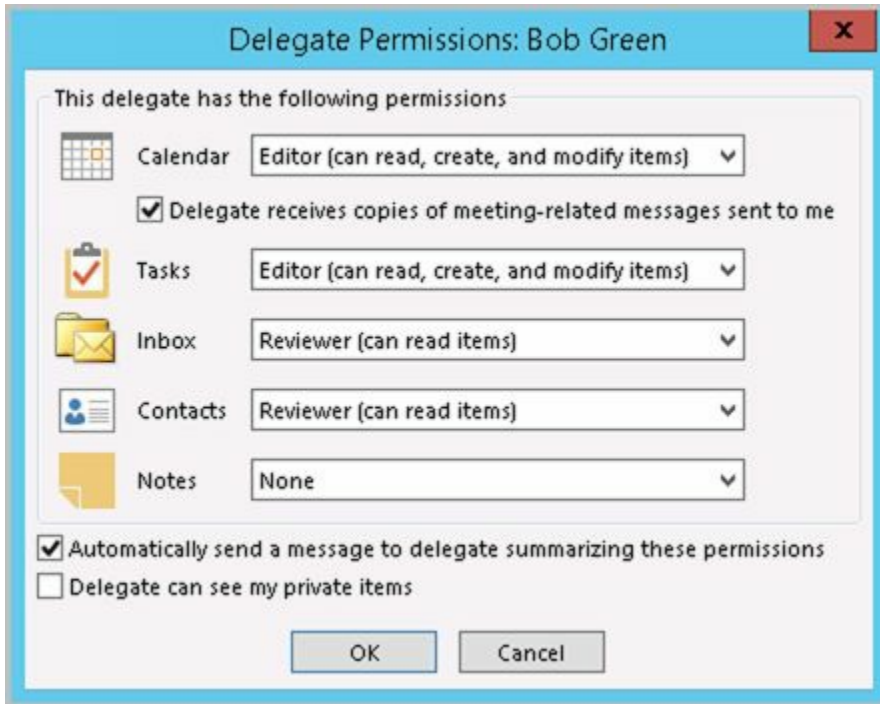
2. The Add Users dialog box appears. To add users, double-click the name of a user who needs access to the mailbox. Repeat this step as necessary for other users, and then click **OK** when you're finished.
3. In the Delegate Permissions dialog box, assign permissions to the delegates for the Calendar, Tasks, Inbox, Contacts, and Notes. The available permissions include

- **None** No permissions
- **Reviewer** Grants read permission only
- **Author** Grants read and create permissions
- **Editor** Grants read, create, and modify permissions

NOTE If the user needs total control over the mailbox, you should grant the user

Editor permission for all items.

4. Click **OK** twice. These changes go into effect when the user restarts Outlook.

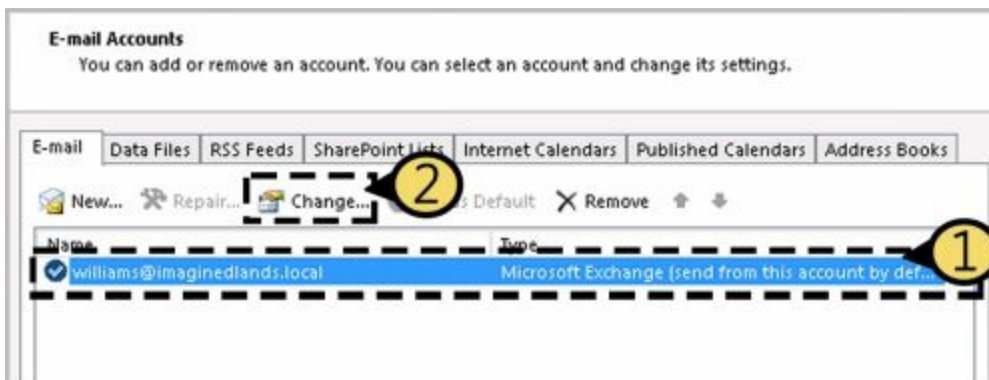


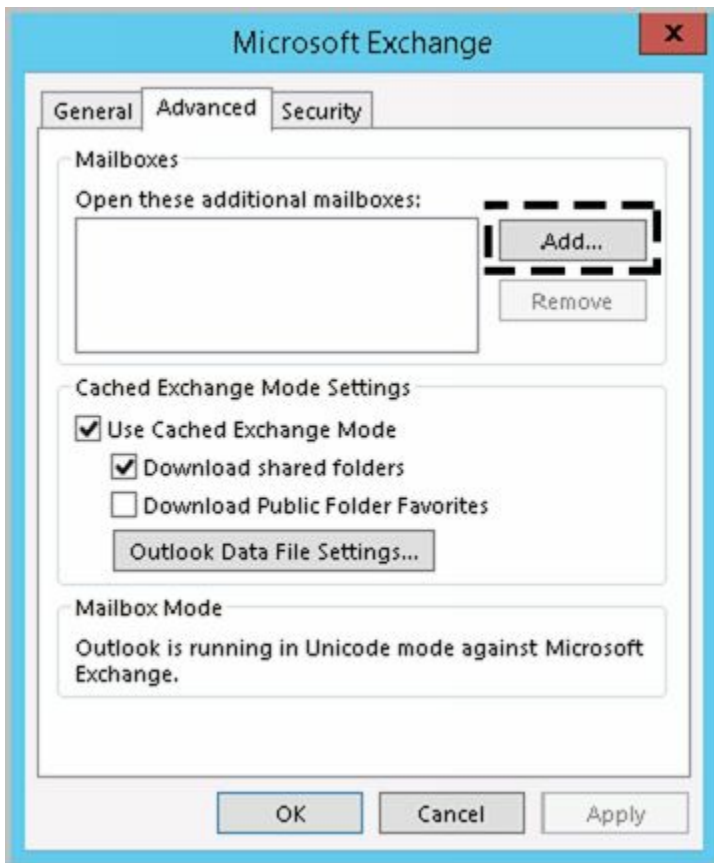
Delegated users can access the mailbox and send mail on behalf of the mailbox owner. To change this behavior, set folder permissions as described later in the “Granting permission to access folders without delegating access” section.

Opening Additional Exchange Mailboxes

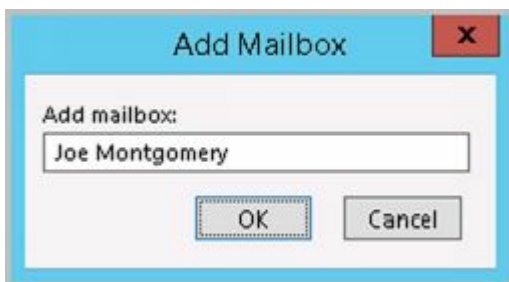
The final step is to let Exchange Server know about the additional mailboxes the user can open. To do this, follow these steps:

1. Have the user who will be accessing additional mailboxes log on and start Outlook.
2. In Outlook 2010, click the Office button, click **Account Settings**, and then select the **Account Settings** option. In Outlook 2013 or Outlook 2016, on the File pane, click **Account Settings**, and then select the **Account Settings** option.
3. Select the Microsoft Exchange Server account, and then click **Change**.





4. In the Change Account dialog box, click **More Settings**.
5. In the Microsoft Exchange dialog box, on the Advanced tab, click **Add** .
6. Type the name of a mailbox to open. Generally, this is the same name as the mail alias for the user or account associated with the mailbox. Click **OK** . Repeat this step to add other mailboxes.



7. Click **Next** , and then click **Finish** .
8. Click **Close** . The additional mailboxes are displayed in the Outlook folder list.

Granting Permission to Access Folders Without Delegating Access

When a mailbox is stored on the server, you can grant access to individual folders in the mailbox. Granting access in this way allows users to add the mailbox to their mail profiles and work with the folder. Users can perform tasks only for which you've granted permission.

To grant access to folders individually, follow these steps:

1. Right-click the folder for which you want to grant access, and then

select **Properties** . In the Properties dialog box, select the **Permissions** tab, as shown in Figure 10-8.

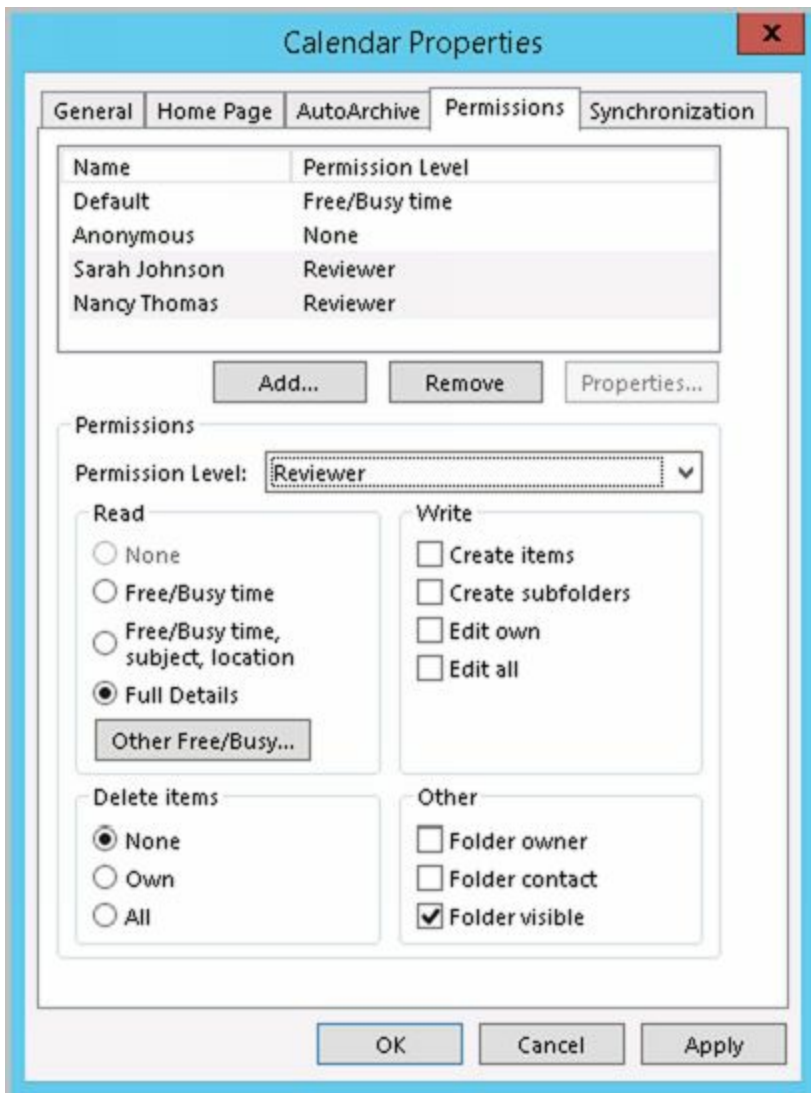
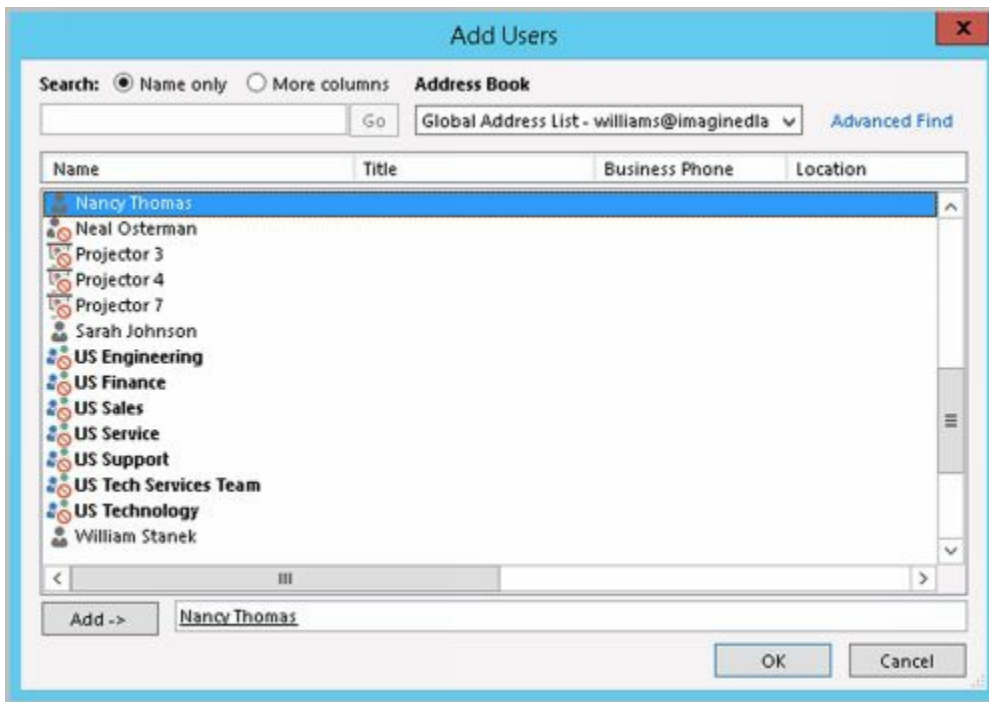


FIGURE 10-8 Granting access to a folder through the Permissions tab.

2. The Name and Permission Level lists display account names and their permissions on the folder. Two special names might be listed:
 - **Default** Provides default permissions for all users.
 - **Anonymous** Provides permissions for anonymous users, such as those who anonymously access a published public folder through the web.
3. To grant permission that differs from the default permission, click **Add** .
4. In the Add Users dialog box, double-click the name of a user who needs access to the mailbox. Repeat this step as necessary for other users, and click **OK** when finished.
5. In the Name and Role lists, select one or more users whose permissions you want to modify. Then use the Permission Level list to assign permissions or select individual permission items. The roles are defined as follows:



- **Owner** Grants all permissions in the folder. Users with this role can create, read, modify, and delete all items in the folder. They can create subfolders and change permissions on folders as well.
- **Publishing Editor** Grants permission to create, read, modify, and delete all items in the folder. Users with this role can create subfolders as well.
- **Editor** Grants permission to create, read, modify, and delete all items in the folder.
- **Publishing Author** Grants permission to create and read items in the folder, to modify and delete items the user created, and to create subfolders.
- **Author** Grants permission to create and read items in the folder and to modify and delete items the user created.
- **Nonediting Author** Grants permission to create and read items in the folder.
- **Reviewer** Grants read-only permission.
- **Contributor** Grants permission to create items but not to view the contents of the folder.
- **None** Grants no permission in the folder.

6. When you're finished granting permissions, click **OK**.

Using Mail Profiles to Customize the Mail Environment

The mail profile used with Outlook determines which information services are available and how they are configured. A default mail profile is created when you install and configure Outlook for the first time. This mail profile is usually called Outlook.

The active mail profile defines the mail setup for the user who is logged on to the computer. You can define additional profiles for the user as well. You can use these additional profiles to customize the user's mail environment for different situations. Here are two scenarios:

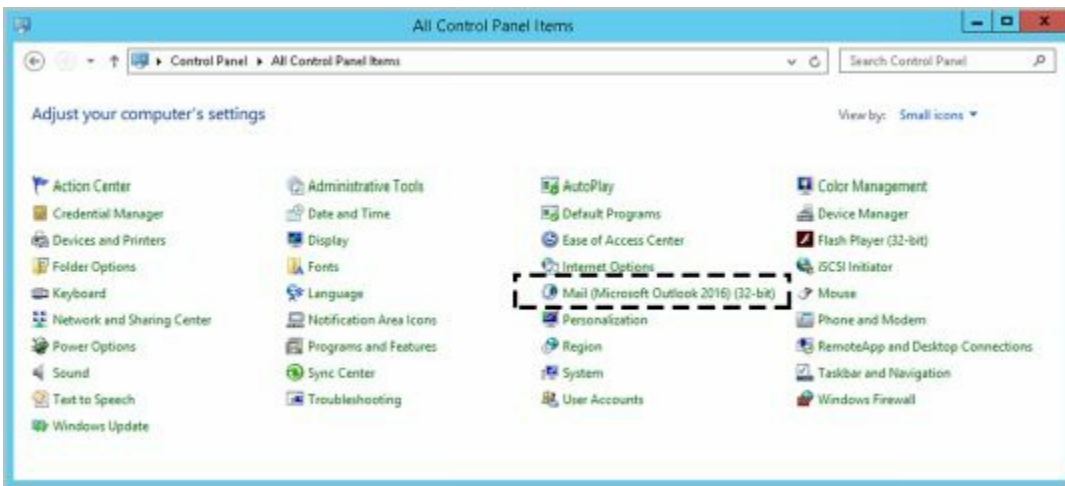
- A manager needs to check the Technical Support and Customer Support mailboxes only on Mondays when she writes summary reports. On other days, the manager doesn't want to see these mailboxes. To solve this problem, you create two mail profiles: Support and Standard. The Support profile displays the manager's mailbox as well as the Technical Support and Customer Support mailboxes. The Standard profile displays only the manager's mailbox. The manager can then switch between these mail profiles as necessary.
- A laptop user wants to check Exchange mail directly while connected to the LAN. When at home, the user wants to use remote mail with scheduled connections. On business trips, the user wants to use SMTP and POP3. To solve this problem, you create three mail profiles: On-Site, Off-Site, and Home. The On-Site profile uses the Exchange Server service with a standard configuration. The Off-Site profile configures Exchange Server for remote mail and scheduled connections. The Home profile uses the Internet mail service instead of the Exchange information service.

Common tasks you'll perform to manage mail profiles are examined in this section.

Creating, Copying, and Removing Mail Profiles

You manage mail profiles through the Mail utility. To access this utility and manage profiles, follow these steps:

1. Exit Outlook. Start the Mail utility. In Control Panel, click **Small Icons** on the View By list and then start the Mail app by clicking its icon or by double-clicking its icon.



2. In the Mail Setup–Outlook dialog box, click **Show Profiles** .
3. As Figure 10-9 shows, you should see a list of mail profiles for the current user.

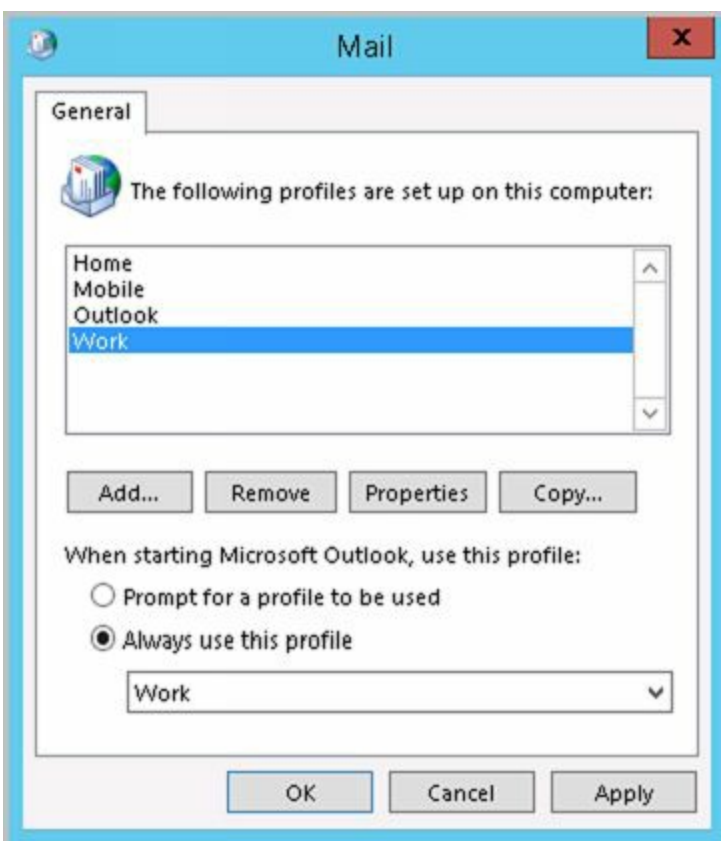


FIGURE 10-9 Using the Mail dialog box to add, remove, or edit mail profiles.

4. Mail profiles for other users aren't displayed. You can now perform the following actions:
 - Click Add to create a new mail profile using the Account Settings Wizard.
 - Delete a profile by selecting it and clicking Remove.
 - Copy an existing profile by selecting it and clicking Copy.
 - View a profile by selecting it and clicking Properties.

Selecting a Specific Profile to use on Startup

You can configure Outlook to use a specific profile on startup or to prompt for a profile to use.

To start with a specific profile, follow these steps:

1. Start the Mail utility. Press the Windows key +I and then click **Control Panel** . In Control Panel, click **Small Icons** on the View By list and then start the Mail app by clicking its icon or by double-clicking its icon.
2. In the Mail Setup–Outlook dialog box, click **Show Profiles** .
3. Select **Always Use This Profile** , and then use the drop-down list to choose the startup profile. Click **OK** .



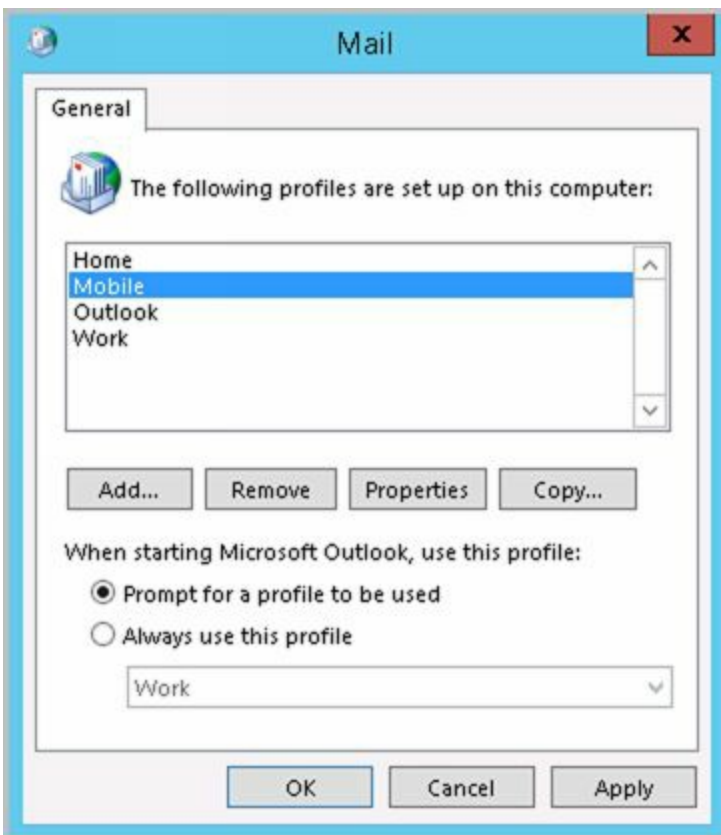
To prompt for a profile before starting Outlook, follow these steps:

1. Start the Mail utility. Press the Windows key +I and then click **Control Panel** . In Control Panel, click **Small Icons** on the View By list and then start the Mail app by clicking its icon or by double-clicking its icon.

@techjob



2. In the Mail Setup–Outlook dialog box, click **Show Profiles** .
3. Select **Prompt For A Profile To Be Used** , and then click **OK** .



Chapter 11. Customizing & Troubleshooting the Exchange Shell

As discussed earlier in the text, the Exchange Management Shell is a command-line management interface built on Windows PowerShell. You use the Exchange Management Shell to manage any aspect of an Exchange Server 2016 configuration that you can manage in the Exchange Admin Center. This means that you can typically use either tool to configure Exchange Server 2016. However, only the Exchange Management Shell has the full complement of available commands, and this means that some tasks can be performed only at the shell prompt.

Running and using the Exchange Management Shell

After you've installed the Exchange management tools on a computer, you can start to use the Exchange Management Shell and the following techniques:

- With Windows 8.1, Windows 10 as well as Windows Server 2012 and Windows Server 2016, you can start Exchange Management Shell by using the Apps Search box. Type **shell** in the Apps Search box, and then select Exchange Management Shell. Or click Start and then choose Exchange Management Shell.
- With Windows 7 and Windows Server 2008 R2, you can start Exchange Management Shell by clicking Start, pointing to All Programs, clicking Microsoft Exchange Server 2016, and then clicking Exchange Management Shell.

The Exchange Management Shell is designed to be run only on domain-joined computers. Whether you are logged on locally to an Exchange server or working remotely, this opens a custom Windows PowerShell console. The console does the following:

1. Connects to the closest Exchange 2016 server using Windows Remote Management (WinRM).
2. Performs authentication checks that validate your access to the Exchange 2016 server and determine the Exchange role groups and roles your account is a member of. You must be a member of at least one management role.
3. Creates a remote session with the Exchange 2016 server. A remote session is a runspace that establishes a common working environment for executing commands on remote computers.

Selecting the shell in this way starts the Exchange Management Shell using your user credentials. This enables you to perform any administrative tasks allowed for your user account and in accordance with the Exchange role groups and management roles you're assigned. As a result, you don't need to run the Exchange Management Shell in elevated, administrator mode, but you can. To do so, right-click Exchange Management Shell, and then click Run As Administrator.

The actual command that runs when you start the shell is:

```
C:\Windows\System32\WindowsPowerShell\v1.0\powershell.exe -noexit -command  
". 'C:\Program Files\Microsoft\Exchange Server\V15\bin\RemoteExchange.ps1';  
Connect-ExchangeServer -auto -ClientApplication:ManagementShell "
```

As you can see, the command starts PowerShell, runs the RemoteExchange.ps1 profile file, and then uses the command Connect-ExchangeServer to establish the remote session. Note the parameters passed in for Connect-ExchangeServer. The –ClientApplication parameter specifies that client-side application is the Exchange Management Shell. The –Auto parameter tells the cmdlet to automatically discover and

try to connect to an appropriate Exchange 2016 server. Discovery works like this:

1. When you run the command on an Exchange 2016 server, the local server is tried first.
2. Next, the command tries to connect to a Mailbox server in the current Active Directory site.
3. Finally, the command tries to connect to a Mailbox server in the current Active Directory site.
4. If no server is available, the command exits.

The RemoteExchange.ps1 profile file sets aliases, initializes Exchange global variables, and loads .NET assemblies for Exchange. It also modifies the standard PowerShell prompt so that it is scoped to the entire Active Directory forest and defines the following Exchange-specific functions:

- **Functions** Allows you to list all available functions by typing **functions** .
- **Get-Exbanner** Displays the Exchange Management Shell startup banner whenever you type **get-exbanner** .
- **Get-Exblog** Opens Internet Explorer and accesses the Exchange blog at Microsoft whenever you type **get-exblog** .
- **Get-Excommand** Allows you to list all available Exchange commands by typing **get-excommand** .
- **Get-Pscommand** Allows you to list all available PowerShell commands by typing **get-pscommand** .
- **Get-Tip** Displays the tip of the day whenever you type **get-tip** .
- **Quickref** Opens Internet Explorer and allows you to download the Exchange Management Shell quick start guide whenever you type **quickref** .

The RemoteExchange.ps1 profile loads the ConnectFunctions.ps1 script, which defines a number of functions that enable AutoDiscover and Connect features. The functions include the following:

- Connect-ExchangeServer
- CreateOrGetExchangeSession
- Discover-EcpVirtualDirectoryForEmc
- Discover-ExchangeServer
- _AutoDiscoverAndConnect
- _CheckServicesStarted
- _ConnectToAnyServer
- _GetCAFEServers
- _GetMailservers
- _GetCurrentVersionServers
- _GetExchangeServersInSite
- _GetHostFqdn
- _GetHubMailboxUMServers
- _GetLocalForest

- [_GetServerFqdnFromNetworkAddress](#)
- [_GetSites](#)
- [_GetWebServiceServers](#)
- [_GetURL](#)
- [_NewExchangeRunSpace](#)
- [_OpenExchangeRunSpace](#)
- [_PrintUsageAndQuit](#)
- [_SelectVdir](#)

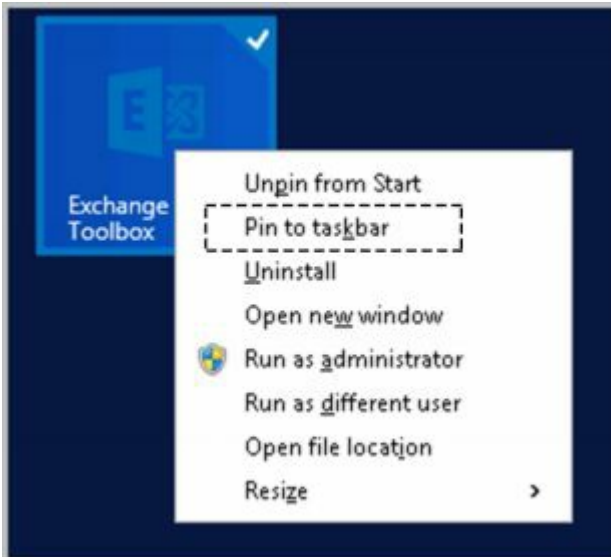
These functions are available for you to use whenever you work with the Exchange Management Shell or have loaded the `ConnectFunctions.ps1` script. However, only `Connect-ExchangeServer`, `CreateOrGetExchangeSession`, `Discover-EcpVirtualDirectoryForEmc` and `Discover-ExchangeServer` are meant to be called directly. The other functions are helper functions.

When you are working with the Exchange Management Shell or have run `ConnectFunctions.ps1`, you can view the source for a function by typing **functions** followed by the name of the function, such as **functions connect-exchangeserver** .

If you want to access Exchange features from a manual remote shell (as discussed later in this chapter under “Using a Manual Remote Shell to Work with Exchange”) or within scripts, you need to load the `RemoteExchange.ps1` profile file. You can find an example of the command required to do this by viewing the properties of the shortcut for the Exchange Management Shell. In the Properties dialog box, the Target text is selected by default. Press `Ctrl+C` to copy this text so that you can use it. For example, if you copy the Target text and paste it into an elevated command prompt (`cmd.exe`), you can access the Exchange Management Shell and work with Exchange Server. If you copy the Target text and paste it into a script, you can be sure that the manual remote session is established when you run the script.

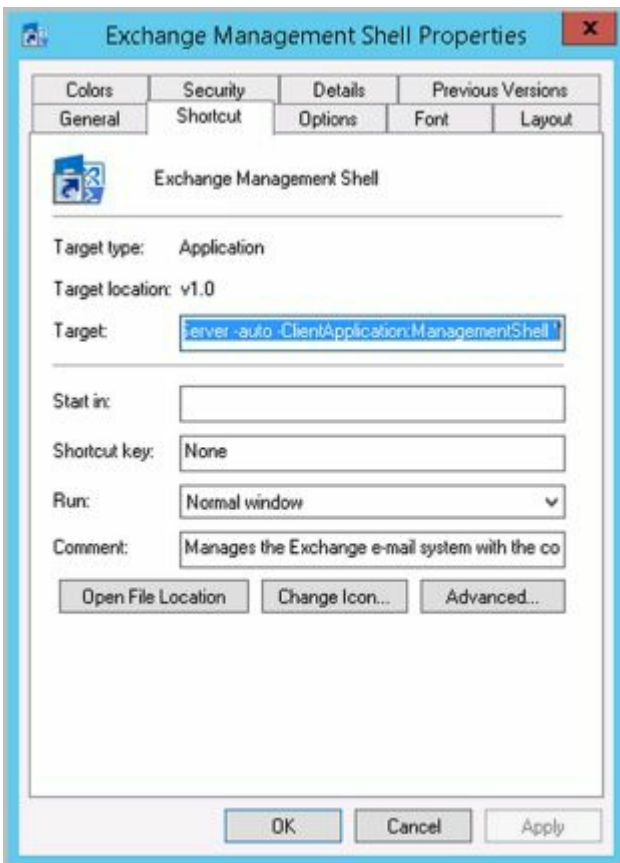
You also can customize the way the Exchange Management Shell is initialized by editing the shortcut properties or by copying the shortcut that starts the Exchange Management shell and then editing the properties. With Windows 8.1, Windows 10 as well as Windows Server 2012 and Windows Server 2016, one way to create a new shortcut for Exchange Management Shell is to do the following:

1. If an Exchange Management Shell shortcut is not pinned to the desktop taskbar, open the Start screen. Next, right-click **Exchange Management Shell** and then select **Pin To Taskbar** .

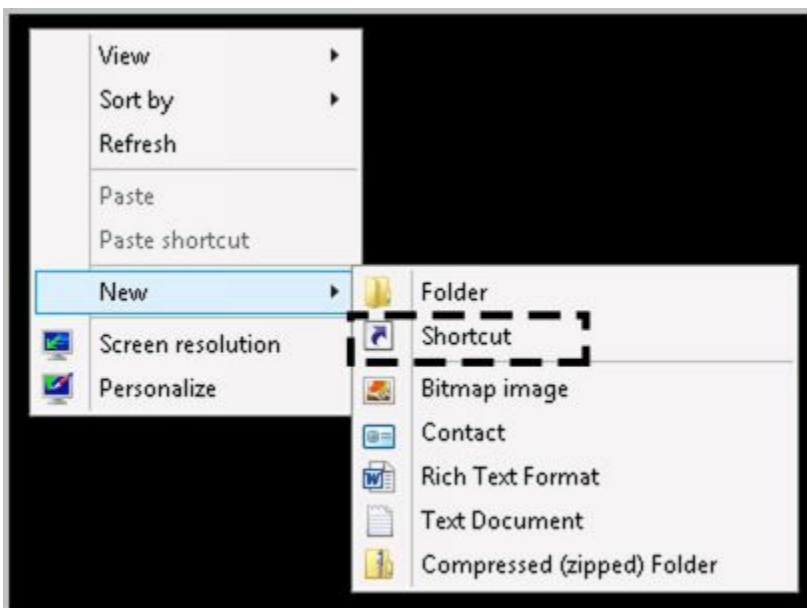


2. On the desktop, right-click the taskbar shortcut for Exchange Management Shell. This displays the Tasks dialog box.
3. In the Tasks dialog box, right-click **Exchange Management Shell** and then select **Properties** . This opens the Properties dialog box for the shortcut with the Shortcut tab selected.

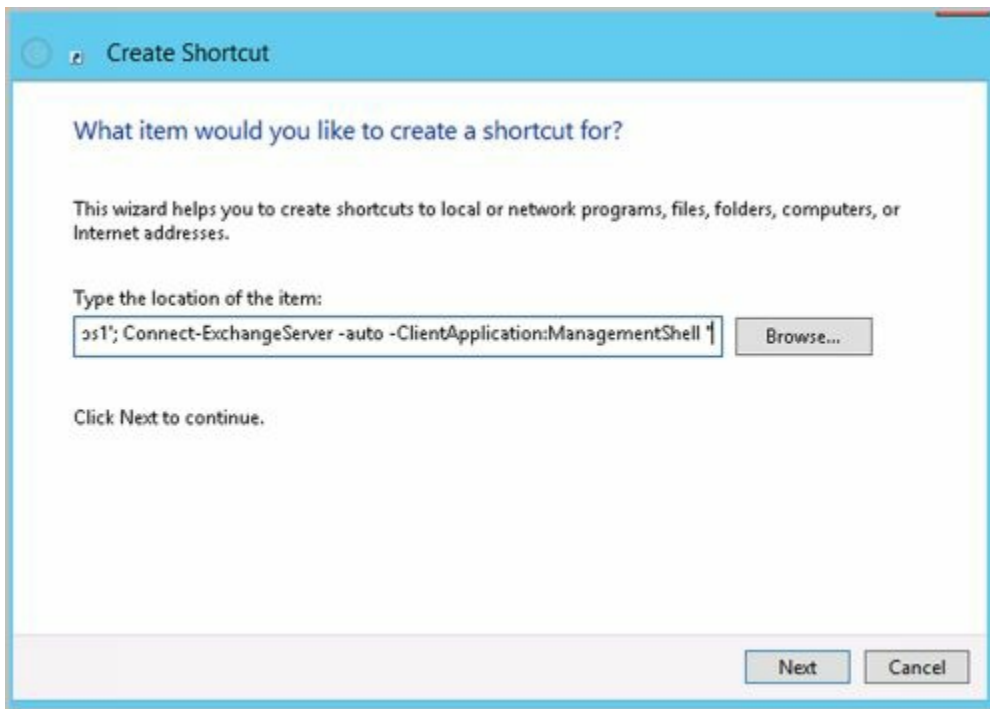




4. Click in the Target box. The text should be selected automatically so you can copy it in the next step. If the text isn't selected, press **Ctrl+A** to select all of the related text.
5. Press **Ctrl+C** to copy the selected text and then click **OK** to close the Properties dialog box.
6. Right-click an open area of the desktop, select **New**, and then select **Shortcut**. This opens the Create Shortcut dialog box.



7. In the Create Shortcut dialog box, click in the Type The Location Of The Item box and then press **Ctrl+V** to paste the previously selected text.



8. Click **Next** . Type a name for the shortcut, such as Custom EMC. Click **Finish** to create the shortcut.
9. Run the shortcut with your custom options by double-clicking it on the desktop.

An extra command must always be added to the Target text. This additional command is Connect-ExchangeServer, a command enabled when the ConnectFunctions.ps1 script runs. To customize the initialization of remote sessions, other parameters are available:

- **-ClearCache** A troubleshooting option that allows you to clear registry entries and exported modules and then re-create the registry settings and import modules again. After you clear the cache, you can try to connect again using options you need.

```
connect-exchangeserver -clearcache
```

- **-Forest** Allows you to specify a single part name or the fully qualified domain name (FQDN) of the Active Directory forest in which to perform discovery. You must be able to authenticate in the forest. User credentials you provide for the **-Username** parameter are not used for discovery. Use with **-Auto**.

```
connect-exchangeserver -auto -forest ForestName
```

- **-Prompt** Prompts you for the FQDN of the Exchange server to connect to. If you use **-Prompt** with **-Auto**, you are prompted only if PowerShell cannot connect automatically. If you use **-Prompt** with **-ServerFqdn**, you are prompted only if PowerShell cannot connect to the specified server.

```
connect-exchangeserver -auto -prompt
```

- **-ServerFqdn** Allows you to specify the FQDN of the Exchange server to connect to.

```
connect-exchangeserver -serverfqdn ExServerFQDN
```

- **-Username** Allows you to specify the user name to use for authentication. You will

be prompted for the user's password. You can also pass in a Credential object. Use with `-ServerFqdn` or `-Auto`.

```
connect-exchangeserver -serverfqdn ExServerFQDN  
-username UserName
```

REAL WORLD When you are working with some cmdlets and objects in PowerShell, you might need to specify a credential for authentication. To do this, use `Get-Credential` to obtain a Credential object and save the result in a variable for later use. Consider the following example:

```
$cred = get-credential
```

PowerShell reads this command, prompts you for a user name and password, and then stores the credentials provided in the `$cred` variable. You also can specify that you want the credentials for a specific user in a specific domain. The following example requests the credentials for the `ExAdmin` account in the `Imaginedlands.com` domain:

```
$cred = get-credential -credential imaginedlands\exadmin
```

A Credential object has `UserName` and `Password` properties that you can work with. Although the user name is stored as a regular string, the password is stored as a secure, encrypted string. Simply pass in the credential instead of the user name as shown in this example:

```
$cred = get-credential -credential imaginedlands\exadmin  
get-hotfix -credential $cred -computername MailServer22
```

IMPORTANT When you prompt for credentials, integrated Windows authentication is used for authentication. However, if the credentials are not set when prompted, such as when the user selects `Cancel`, Kerberos authentication is used with the user's default credentials.

REAL WORLD Where a domain name is required for credentials, you typically can use either the NET BIOS domain name or the DNS domain name. In the previous examples, I entered the NET BIOS domain name **pocket-consulta** rather than the DNS name **imaginedlands.com**.

When you call `Connect-ExchangeServer`, the function does one of two things: It opens a remote session by using implicit credentials (the credentials of the user who is running Exchange Management Shell) or by using specified credentials (credentials you've explicitly provided). One of the final things `Connect-ExchangeServer` does is call `_OpenExchangeRunspace`, which in turn calls `_NewExchangeRunspace` to establish the remote session.

In the script, the core code for `_OpenExchangeRunspace` is:

```
$global:remoteSession = _NewExchangeRunspace $fqdn $credential $UseWIA $SuppressError  
$ClientApplication $AllowRedirection
```

And the core code for `_NewExchangeRunspace` is:

```
$so = New-PSSessionOption -OperationTimeout $sessionOptionsTimeout  
-IdleTimeout $sessionOptionsTimeout -OpenTimeout $sessionOptionsTimeout;
```

```
New-PSSession -ConnectionURI "$connectionUri" -ConfigurationName  
Microsoft.Exchange -SessionOption $so
```

The code sample creates a global variable named `$remoteSession` to hold the remote session. A global variable is used to ensure that the session remains active and available when the script exits. The session is established using `New-PSSession` with a connection URI for a particular Exchange server. For example, if the Exchange server's FQDN is `MailServer15.Imaginedlands.local`, the connection URI is `https://mailserver15.imaginedlands.local/powershell`. The `-ConfigurationName` parameter sets the configuration namespace as `Microsoft.Exchange` (in place of the default `Microsoft.PowerShell`). The `-SessionOption` parameter sets session options that were defined previously using the `New-PSSessionOption` cmdlet. The session options include the operation timeout value, the idle timeout value, and the open session timeout value. By default, all three are set to 180,000 milliseconds (180 seconds) via the `$sessionOptionsTimeout` variable defined in the first section of the `ConnectFunctions.ps1` script.

You can use the `MsExchEmsTimeout` environment variable to set the default timeout values. If you set this environment variable to a value of 900,000 milliseconds or less (15 minutes or less), the timeouts are set accordingly. If you set this environment variable to a value greater than 900,000 milliseconds, the timeout values revert to the 3-minute default value.

REAL WORLD When you connect to Exchange Admin Center in a browser, the browser version determines your experience level, and the location of your mailbox determines whether you see the console for Exchange 2010, Exchange 2013, or Exchange 2016. This is not the case when you are working with the shell. With the shell, the experience level is always set to FULL. Further, the `HKLM:\SOFTWARE\Microsoft\ExchangeServer\v15\Setup` key in the registry is examined to determine the Exchange version and the build number, and then this information is used to set the client version compatibility level. Thus, a precise connection URI is set as `http://$fqdn/powershell?serializationLevel=Full;ExchClientVer=$clientVersion`.

Managing the PowerShell Application

Microsoft Internet Information Services (IIS) handles every incoming request to a website within the context of a web application. A web application is a software program that delivers web content to users over HTTP or HTTPS. Each website has a default web application and one or more additional web applications associated with it. The default web application handles incoming requests that aren't assigned to other web applications. Additional web applications handle incoming requests that specifically

reference a particular application.

When you connect to a server using a URL, such as `https://mailserver15.imaginedlands.local/powershell`, you are performing remote operations via the PowerShell application running on the web server providing Exchange services. Like all web applications, the PowerShell application has a virtual directory associated with it. The virtual directory sets the application name and maps the application to the physical directory that contains the application's content.

You can manage the PowerShell application using IIS Manager GUI and the Exchange Management Shell. The related commands for the Exchange Management Shell are:

- **Get-PowerShellVirtualDirectory** Displays information about the PowerShell application running on the web server providing services for Exchange.

```
Get-PowerShellVirtualDirectory [-Identity 'AppName']  
[-DomainController 'DomainControllerName']
```

```
Get-PowerShellVirtualDirectory -Server 'ExchangeServerName'  
[-DomainController 'DomainControllerName']
```

- **New-PowerShellVirtualDirectory** Creates a new PowerShell application running on the web server providing services for Exchange.

```
New-PowerShellVirtualDirectory -Name 'AppName'  
[-AppPoolId 'AppPoolName'] [-BasicAuthentication <$true | $false>]  
[-CertificateAuthentication <$true | $false>] [-DomainController  
'DomainControllerName'] [-ExternalUrl 'URL'] [-InternalUrl 'URL']  
[-Path 'PhysicalDirectoryPath']  
[-WindowsAuthentication <$true | $false>]
```

- **Remove-PowerShellVirtualDirectory** Removes a specified PowerShell application running on the web server providing services for Exchange.

```
Remove-PowerShellVirtualDirectory -Identity 'AppName'  
[-DomainController 'DomainControllerName']
```

- **Set-PowerShellVirtualDirectory** Modifies the configuration settings for a specified PowerShell application running on the web server providing services for Exchange.

```
Set-PowerShellVirtualDirectory -Identity 'AppName'  
[-BasicAuthentication <$true | $false>] [-CertificateAuthentication  
<$true | $false>] [-DomainController 'DomainControllerName']  
[-ExternalUrl 'URL'] [-InternalUrl 'URL']  
[-LiveIdBasicAuthentication <$true | $false>]  
[-WindowsAuthentication <$true | $false>]
```

At the Exchange Management Shell prompt, you can confirm the location of the PowerShell application by typing `get-powershellvirtualdirectory`.

`GetPowerShellVirtualDirectory` lists the name of the application, the associated

directory and website, and the server on which the application is running, as shown in the following example:

Name	Server
-----	-----
PowerShell (Default Web Site)	CorpServer45

In this example, a standard configuration is being used where the application named *PowerShell* is running on Default Web Site on CorpServer45. You can use `Set-PowerShellVirtualDirectory` to specify the internal and external URL to use as well as the permitted authentication types. Authentication types you can enable or disable include basic authentication, Windows authentication, certificate authentication, and Live ID basic authentication. You can use `New-PowerShellVirtualDirectory` to create a new PowerShell application on the web server providing services for Exchange and `Remove-PowerShellVirtualDirectory` to remove a PowerShell application.

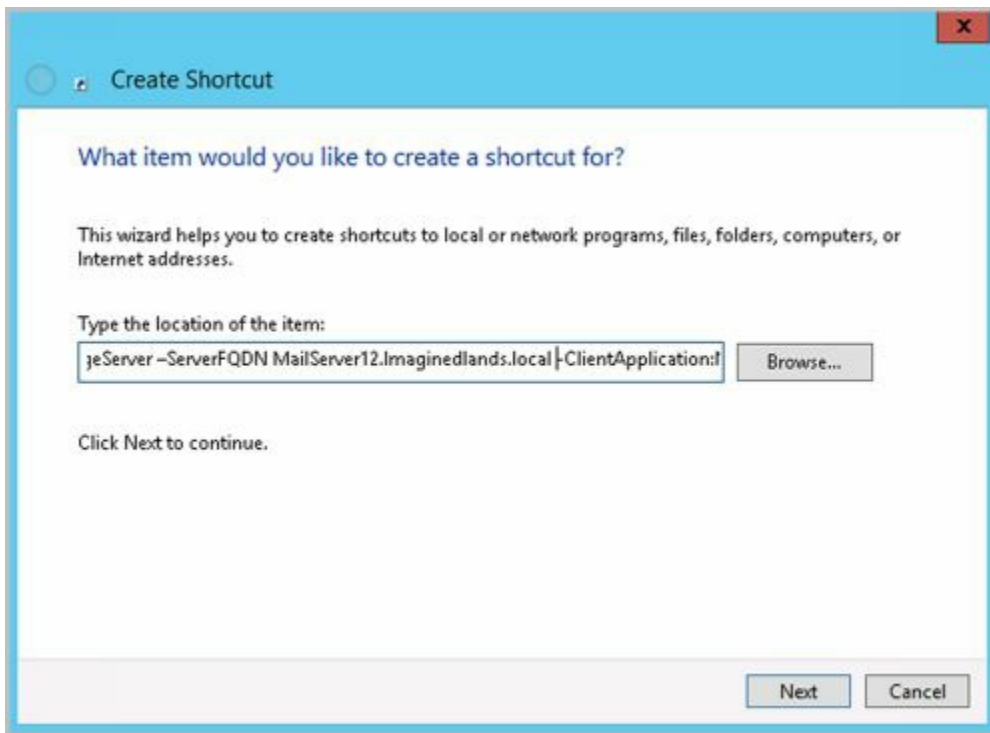
REAL WORLD Any change you make to the PowerShell virtual directory configuration requires careful pre-planning. For every potential change, you'll need to determine whether you need to modify the WinRM configuration and the PowerShell path in `ConnectFunctions.ps1` scripts on management computers and Exchange servers as well as the specific changes you'll need to make with regard to IIS on your Mailbox servers.

Microsoft cautions against modifying the default configuration for the PowerShell virtual directory as any mistakes you make could prevent you from managing Exchange Server. Because Exchange configuration data is stored in Active Directory and the affected IIS metabase, you would need to be able to restore Exchange data in Active Directory and the affected IIS metabase to a previous state to recover.

Customizing Exchange Management Shell

Now that you know how the Exchange Management Shell environment works, you can more easily customize the shell to work the way you want it to. One way to do this is to modify the menu shortcut that starts the Exchange Management Shell or create copies of this menu shortcut to change the way the Exchange Management Shell starts. For example, if you want to connect to a named Exchange server rather than any available Exchange server, you can do the following:

1. In the Properties dialog box for the shortcut that starts Exchange Management Shell, the Target text is selected by default. Press the right arrow key to move to the end of the command text.
2. Delete `-Auto` and type `-ServerFqdn` followed by the FQDN of the Exchange server, such as `-ServerFQDN MailServer12.Imaginedlands.local`. Click **OK**.



That said, this entire sequence of tasks is meant to simplify the task of establishing an interactive remote session with a single Exchange server. As implemented in the default configuration, you have a one-to-one, interactive approach for remote management, meaning you establish a session with a specific remote server and work with that specific server simply by executing commands.

When you are working with PowerShell outside of Exchange Management Shell, you might want to use the `Enter-PSSession` cmdlet to start an interactive session with an Exchange server or any other remote computer. The basic syntax is **Enter-PSSession *ComputerName***, where *ComputerName* is the name of the remote computer, such as the following:

```
enter-pssession mailserver15
```

After you enter this command, the command prompt changes to show that you are connected to the remote computer, as shown in the following example:

```
[MailServer15]: PS C:\Users\wrstaneck.cpandl\Documents>
```

Now, the commands that you type run on the remote computer just as if you had typed them directly on the remote computer. In most cases, you need to ensure you are running an elevated, administrator shell and that you pass credentials along in the session. When you connect to a server in this way, you use the standard PowerShell remoting configuration and do not go through the PowerShell application running on a web server. You can end the interactive session by using the command `Exit-PSSession` or typing **exit**.

To access an Exchange server in the same way as the `ConnectFunctions.ps1` script, you need to use the `-ConnectionURI` parameter to specify the connection URI, the `-ConfigurationName` parameter to specify the configuration namespace, the `-Authentication` parameter to set the authentication type to use, and optionally, the `-`

SessionOption parameter to set session options. Consider the following example:

```
enter-psession -connectionURI http://mailserver12.imaginedlands.local/powershell  
-ConfigurationName Microsoft.Exchange -Authentication Kerberos
```

Here, you set the connection URI as

`https://mailserver12.imaginedlands.local/powershell`, set the configuration namespace as `Microsoft.Exchange`, and use Kerberos authentication with the implicit credentials of your user account. If you don't specify the authentication method, the default authentication method for WinRM is used. If you want to use alternate credentials, you can pass in credentials as shown in this example:

```
$cred = get-credential -credential imaginedlands\williams
```

```
enter-psession -connectionURI https://mailserver12.imaginedlands.local/powershell  
-ConfigurationName Microsoft.Exchange -credential $cred  
-Authentication Kerberos
```

Here, you set the connection URI as

`https://mailserver12.imaginedlands.local/powershell`, set the configuration namespace as `Microsoft.Exchange`, and use alternate credentials. When PowerShell reads the `Get-Credential` command, you are prompted for the password for the specified account. Because the authentication type is not defined, the session uses the default authentication method for WinRM.

To put this all together, one way to create a script that runs on an Exchange server is to run the `RemoteExchange.ps1` profile file and then run the `ConnectFunctions.ps1` script to autoconnect to Exchange. The commands you insert into your script to do this are the following:

```
$s = $env:ExchangeInstallPath + "bin\RemoteExchange.ps1"  
&$s  
$t = $env:ExchangeInstallPath + "bin\ConnectFunctions.ps1"  
&$t
```

Here, you define variables that point to the `RemoteExchange.ps1` and `ConnectFunctions` scripts in the Exchange installation path, and then you use the `&` operator to invoke the scripts. The environment variable `ExchangeInstallPath` stores the location of the Exchange installation. If you enter the full path to a script, you don't need to assign the path to a variable and then invoke it. However, you then have a fixed path and might need to edit the path on a particular Exchange server. Be sure to run the script at an elevated, administrator PowerShell prompt.

To create a script that runs on your management computer and then executes commands remotely on an Exchange server, insert commands in your script to create a new session and then invoke commands in the session using the techniques discussed in the next section.

Performing One-to-Many Remote Management

PowerShell also lets you perform one-to-many remote management. To do so, you must work with an elevated, administrator shell and can either invoke remote commands on multiple computers or establish remote sessions with multiple computers. When you remotely invoke commands, PowerShell runs the commands on the remote computers, returns all output from the commands, and establishes connections to the remote computers only for as long as is required to return the output. When you establish remote sessions, you can create persistent connections to the remote computers and then execute commands within the session. Any command you enter while working in the session is executed on all computers to which you are connected, whether this is 1 computer, 10 computers, or 100 computers.

TIP As discussed in Chapter 1, “Welcome to Exchange Server 2016,” WinRM must be appropriately configured on any computer you want to remotely manage. While WinRM is configured on Exchange servers and most others computers running current versions of Windows and Windows Server, WinRM listeners generally are not created by default. You can create the required listeners by running `winrm quickconfig`.

The following command entered as a single line invokes the `Get-Service` and `Get-Process` commands on the named servers:

```
invoke-command -computername MailServer12, Mailserver19, MailServer32  
-scriptblock {get-service; get-process}
```

The following command establishes a remote session with the named computers:

```
$s = new-PSSession -computername MailServer12, Mailserver19, MailServer32  
-Credential Cpandl\WilliamS
```

When you connect to a server in this way, you use the standard PowerShell remoting configuration and are not going through the PowerShell application running on a web server. After you establish the session, you can then use the `$s` session with `Invoke-Command` to return commands on all remote computers you are connected to. This example looks for stopped Exchange services on each computer:

```
invoke-command -session $s  
-scriptblock {get-service mse* | where { $_.status -eq "stopped"}}
```

In this example, you pipe the output of `Get-Service` to the `Where-Object` cmdlet and filter based on the `Status` property. Because the `$_` automatic variable operates on the current object in the pipeline, PowerShell examines the status of each service in turn and lists only those that are stopped in the output.

In addition to working with remote commands and remote sessions, some cmdlets have a `ComputerName` parameter that lets you work with a remote computer without using Windows PowerShell remoting. PowerShell supports remote background jobs as well. A background job is a command that you run asynchronously in an interactive or noninteractive session. When you start a background job, the command prompt returns immediately, and you can continue working while the job runs. For a complete

discussion of these remoting features, see *Windows PowerShell: The Personal Trainer* (Stanek & Associates, 2014).

Using a Manual Remote Shell to Work with Exchange

Although the easiest way to work remotely with Exchange 2016 is to install the management tools on your computer, you can connect to and manage Exchange 2016 if you don't have the management tools installed. To do this, you can use a manual remote shell to connect to an Exchange 2016 server. However, you lose the benefits of the preconfigured tools which set up the environment and manage the Exchange connection for you. You also can use a manual remote shell to connect to and work with Exchange Online.

Preparing to Use the Remote Shell

As you might expect, there are several prerequisites for creating a manual remote shell. The computer you use to connect an Exchange server must be running a current version of Windows or Windows Server.

The computer must have Windows Management Framework, which includes Windows PowerShell and WinRM, and Microsoft .NET Framework. Although current versions of Windows and Windows Server include these components, Windows 7 and Windows Server 2008 R2 do not.

REAL WORLD When you install the Mailbox server role for Exchange 2016, the server is configured automatically with a Windows PowerShell gateway that is configured as a proxy service. This proxy service allows you to run remote commands in web browsers and in remote sessions. Whenever you work with Exchange Admin Center or Exchange Management Shell, the commands are executed via this proxy—even if you logged on locally.

Before you can work remotely, WinRM must be running and the authentication mechanisms you want to use must be enabled. As Exchange Online uses Basic authentication, you may need to enable this. At an elevated, administrator PowerShell prompt, enter the following commands to check the status of WinRM:

```
get-service "winrm"
```

If WinRM isn't running, start the service by entering:

```
start-service "winrm"
```

Next, ensure that the authentication mechanisms you want to use are enabled for use with WinRM. To do this, enter the following command:

```
winrm get winrm/config/client/auth
```

Although you are working in the PowerShell window, this command is passed through to the command prompt and the output states the status of available authentication mechanisms:

Auth

```
Basic = false  
Digest = true  
Kerberos = true  
Negotiate = true  
Certificate = true  
CredSSP = false
```

If Basic authentication isn't enabled and you want to work with Exchange Online, you must enable it. Unfortunately, there's no easy way to pass a complex command through to the command prompt. Because of this, you'll need to open an elevated command prompt and then enter the following command:

```
winrm set winrm/config/client/auth '@{Basic="true"}'
```

IMPORTANT Exchange Management Shell and Exchange Admin Center require integrated Windows authentication. Exchange Online uses Basic authentication.

Once you've ensured WinRM is running and configured appropriately, you can check the status of script execution by entering the following command at the PowerShell prompt:

```
Get-ExecutionPolicy
```

Windows PowerShell script execution must be enabled on your computer. Typically, you'll want to use the RemoteSigned execution policy. If so, enter the following command at an elevated, administrator PowerShell prompt:

```
Set-ExecutionPolicy RemoteSigned
```

When you connect to a remote Exchange server, you can use your current network credentials or you can specify another set of credentials. Either way, the user account that you want to use for remote management must be a member of a management role group or be enabled for remote shell.

By default, when you create a new mailbox user for Exchange 2016 or Exchange Online, the mailbox user has remote PowerShell enabled. You can view the access status for all users in the Exchange organization by entering the following command:

```
Get-User -ResultSize unlimited |  
Format-Table Name,DisplayName,RemotePowerShellEnabled
```

If you want to display a list of only users who don't have access, you could filter the results for this value by running the following command instead:

```
Get-User -ResultSize unlimited -Filter {RemotePowerShellEnabled -eq $false}
```

Set the filtered value to \$true if you want to see a list of only users who have access. You can check the access status of a specific user as well by specifying the SAM account name, display name, or login name of the user, as shown in these examples:

```
Get-User "williams" | Format-List RemotePowerShellEnabled
```

```
Get-User "William Stanek" | Format-List RemotePowerShellEnabled
```

```
Get-User "williams@imaginedlands.com" | Format-List
RemotePowerShellEnabled
```

Connecting Manually to Exchange 2016 Servers

In an elevated, administrator Windows PowerShell window, you can establish a connection to the remote Exchange server using a PowerShell session. When your management computer is joined to the domain, you can use either HTTP or HTTPS with Kerberos authentication to establish the session. However, HTTPS is normally disabled by default in the client configuration. The basic syntax is:

```
$Session = New-PSSession -ConfigurationName Microsoft.Exchange
-ConnectionUri http:// Exchange2016MSName /PowerShell/
-Authentication Kerberos
```

where *Exchange2016MSName* is the host name or FQDN of the Exchange 2016 Mailbox server to which you want to connect and *PowerShell* is the name of the PowerShell virtual directory on the server, such as:

```
$Session = New-PSSession -ConfigurationName Microsoft.Exchange
-ConnectionUri http://mailserver35.imaginedlands.com/PowerShell/
-Authentication Kerberos
```

With Kerberos authentication, your current credentials are used to establish the session. Keep in mind that with Kerberos authentication you must use the server name or the FQDN and cannot use an IP address.

If you want to use an authentication mechanism other than Kerberos or your computer isn't connected to a domain, you must use HTTPS as the transport (or the destination server must be added to the TrustedHosts configuration settings for WinRM, and HTTP must be enabled in the client configuration). You also must explicitly pass in a credential using the `-Credential` parameter.

You also can specify the authentication mechanism, such as Basic, Digest or Negotiate. All communications are encrypted with HTTPS. The modified syntax is then:

```
$Session = New-PSSession -ConfigurationName Microsoft.Exchange
-ConnectionUri https:// Exchange2016MSNameOrIP /PowerShell/
-Authentication Negotiate -Credential Credential
```

where *Exchange2016MSNameOrIP* is the FQDN or IP address of the Exchange 2016 Mailbox server to which you want to connect, *PowerShell* is the name of the PowerShell virtual directory on the server, and *Credential* sets the user name under which the session is established. Consider the following example:

```
$Session = New-PSSession -ConfigurationName Microsoft.Exchange
-ConnectionUri https://mailserver35.imaginedlands.com/PowerShell/
-Authentication Negotiate -Credential imaginedlands\williams
```

Here, you establish a session with MailServer35 using integrated Windows

authentication and store this session in the `$Session` object. As you are passing in a credential for Williams, you are prompted for and must enter the account password. You also can store the credential in a Credential object and then use `Get-Credential` to request the credentials. The syntax then becomes:

```
$Cred = Get-Credential
$Session = New-PSSession -ConfigurationName Microsoft.Exchange
-ConnectionUri https://mailserver35.imaginedlands.com/PowerShell/
-Authentication Negotiate -Credential $Cred
```

IMPORTANT Regardless of whether you use Kerberos or another authentication mechanism, the Exchange server's SSL certificate must contain a common name (CN) that matches the identifier you are using. Otherwise, you won't be able to connect.

Connecting Manually to Exchange Online

Connecting manually to Exchange Online is similar to connecting manually to on-premises servers running Exchange 2016. In an elevated, administrator Windows PowerShell window, you can establish a connection to Exchange Online using a PowerShell session. You can use a stand-alone computer or a domain-joined computer that meets the requirements discussed earlier under "Preparing to use the remote shell."

The basic syntax for connecting manually to Exchange Online is:

```
$Cred = Get-Credential
$Session = New-PSSession -ConfigurationName Microsoft.Exchange
-ConnectionUri https://ps.outlook.com/powershell/
-Authentication Basic -Credential $Cred -AllowRedirection
```

Here, you use HTTPS with Basic authentication for the session and establish a connection to the Exchange Online URL provided by Microsoft, which typically is `https://ps.outlook.com`. To establish the connection, you must pass in your Exchange Online user name and password. This example stores credentials in a Credential object and then uses `Get-Credential` to prompt for the required credentials. You also could specify the credentials explicitly, as shown here:

```
$Session = New-PSSession -ConfigurationName Microsoft.Exchange
-ConnectionUri https://ps.outlook.com/powershell/
-Authentication Basic -Credential wrs@imaginedlands.onmicrosoft.com
-AllowRedirection
```

Here, you are prompted for the password for the account.

NOTE When you work with Exchange Online, keep in mind that not all of the cmdlets are available as compared to an on-premises installation. This is because the operating environments are different. Exchange Online runs on Windows Azure rather than Windows Server. You can connect to and work directly with the Microsoft Online service and Windows Azure as discussed in Chapter 2, "Working with Exchange Online."

Managing Remote Sessions

After you establish a session with an Exchange 2016 server or Exchange Online, you must import the server-side PowerShell session into your client-side session by running the following command:

```
Import-PSSession $Session
```

You can then work with the remote server.

When you are finished, you should disconnect the remote shell from Exchange server. It's important to note that, beginning with Windows PowerShell 3.0, sessions are persistent by default. When you disconnect from a session, any command or scripts that are running in the session continue running, and you can later reconnect to the session to pick up where you left off. You also can reconnect to a session if you were disconnected unintentionally, such as by a temporary network outage.

IMPORTANT With Exchange Online, each account can have only three connections to sever-side sessions at a time. If you close the PowerShell window without disconnecting from the session, the connection remains open for 15 minutes and then disconnects.

To disconnect a session without stopping commands or releasing resources, run the following command:

```
Disconnect-PSSession $Session
```

The \$Session object was instantiated when you created the session. As long as you don't exit the PowerShell window in which this object was created, you can use this object to reconnect to the session by entering:

```
Connect-PSSession $Session
```

When you are completely finished with the session, you should remove it. Removing a session stops any commands or scripts that are running, ends the session, and releases the resources the session was using. Remove a session by running the following command:

```
Remove-PSSession $Session
```

Troubleshooting Exchange Management Shell

Note that the `ConnectionFunctions.ps1` script relies on your organization having a standard Exchange Server configuration. By default, Exchange is configured for management using HTTP with the URL `http://ServerName/powershell`. If you've modified the Web Server configuration on your Exchange servers to use a different path, such as might be required to enhance security, you need to update the connection URIs used in the `ConnectionFunctions.ps1` script.

When you invoke the PowerShell application, the web server to which you connect runs the PowerShell plug-in (`Pwrshplugin.dll`) and the Exchange Authorization plug-in (`Microsoft.Exchange.AuthorizationPlugin.dll`). The PowerShell plug-in runs as a `Microsoft.Exchange` shell and has the following initialization parameters:

- `PSVersion`, which sets the PowerShell version as appropriate
- `ApplicationBase`, which sets the base path for the Exchange server as `%ExchangeInstallPath%Bin`
- `AssemblyName`, which sets the name of the .NET assembly to load as `Microsoft.Exchange.Configuration.ObjectModel.dll`

The Authorization plug-in handles Exchange authorization and authentication. Together, these plug-ins create an authorized shell environment for the remote session.

The physical directory for the PowerShell application is `%ExchangeInstallPath%\ClientAccess\PowerShell`. This application runs in the context of an application pool named `MSExchangePowerShellAppPool`. In a large organization, you might want to optimize settings for this and other application pools, as discussed in the *IIS Web Applications, Security & Maintenance: The Personal Trainer* (Stanek & Associates, 2015).

In the `%ExchangeInstallPath%\ClientAccess\PowerShell` directory on your server, you'll find a `web.config` file that defines the settings for the PowerShell application. This file contains a role-based access control (RBAC) configuration section that loads the assemblies and web controls for the application.

TIP Microsoft recommends against changing the PowerShell application configuration. However, there's nothing magical or mystical about the PowerShell application or `MSExchangePowerShellAppPool`. You can re-create these features to enable remote management in alternate configurations, such as on nondefault websites or websites with alternate names. However, be sure to copy the PowerShell application's `web.config` file to the physical directory for your base application. Before you make any changes to a live production environment, you should plan and test your changes in a nonproduction test environment.

The web server to which you connect processes your remote actions via the Exchange Control Panel (ECP) application running on the default website. With Exchange 2016,

you see the ECP as the Exchange Admin Center. The physical directory for this application is %ExchangeInstallPath%\ClientAccess\Ecp. This application runs in the context of an application pool named MSEExchangeECPAppPool.

In the %ExchangeInstallPath%\ClientAccess\ECP directory on your server, you'll find a web.config file that defines the settings for the ECP application. This file contains an RBAC configuration section that loads the assemblies and web controls for the application.

Because of the interdependencies created by accessing Exchange via web applications, you'll want to examine related features as part of troubleshooting any issues you experience with remote sessions. Generally, your troubleshooting should follow these steps:

1. Examine the status and configuration of the WinRM on your local computer and the target Exchange server. The service must be started and responding.
2. Check the settings of any firewall running on your local computer, the target Exchange server, or any device between the two, such as a router with a firewall.
3. Check the status of the World Wide Web Publishing Service on the Exchange server. The service must be started and responding.
4. Check the configuration settings of the PowerShell and ECP applications on the web server. By default, the applications don't have access restrictions, but another administrator could have set restrictions.
5. Check the status of MSEExchangePowerShellAppPool and MSEExchangeECPAppPool. You might want to recycle the application pools to stop and then start them.
6. Check the configuration settings of MSEExchangePowerShellAppPool and MSEExchangeECPAppPool. By default, the application pools are configured to use only one worker process to service requests.
7. Check to ensure the PowerShell application's web.config file is present in the physical directory for the application, and also that the file has the appropriate settings.
8. Check to ensure the ECP application's web.config file is present in the physical directory for the application and also that the file has the appropriate settings.

Chapter 12. Customizing & Configuring Exchange Security

You manage Exchange security using either the Active Directory tools or the Exchange management tools. In Active Directory, you manage security using permissions. Users, contacts, and security groups all have permissions assigned to them. These permissions control the resources that users, contacts, and groups can access and the actions they can perform. You use auditing to track the use of these permissions, as well as log ons and log offs. In addition to the standard permissions, Exchange also supports *role-based access control* (RBAC), which are unique to Exchange.

Configuring Standard Exchange Permissions

Active Directory is the central repository for information in domains. Because Active Directory also stores most Exchange information, you can use the features of Active Directory to manage standard permissions for Exchange across the organization.

Assigning Permissions: Exchange Server and Online

Users, contacts, and security groups are represented in Active Directory as objects. These objects have many attributes that determine how they are used. The most important attributes are the permissions assigned to the objects. Permissions grant or deny access to objects and resources. For example, you can grant a user the right to create public folders but deny that same user the right to create mail-enabled contacts.

Permissions assigned to an object can be applied directly to the object, or they can be inherited from another object. Generally, objects inherit permissions from *parent objects*. A parent object is an object that is above another object in the object hierarchy. However, you can override inheritance. One way to do this is to assign permissions directly to an object. Another way is to specify that an object shouldn't inherit permissions.

In Exchange Server 2016, permissions are inherited through the organizational hierarchy. The root of the hierarchy is the *domain*. All other containers in the tree inherit the permissions of the domain container. Sometimes, however, you want to create structures that represent parts of the organization or you want to limit administrative access for part of the organization. To do this, you use organizational units. *Organizational units* (OUs) are containers for objects that you not only want to group together but that you also want to manage together.

For the management of Exchange information and servers, Exchange Server 2016 uses several predefined groups. These predefined security groups have permissions to manage the Exchange organization, Exchange servers, and Exchange recipient data in Active Directory. In Active Directory Users And Computers, you can view and work with the Exchange-related groups using the Microsoft Exchange Security Groups organizational unit (see Figure 12-1).

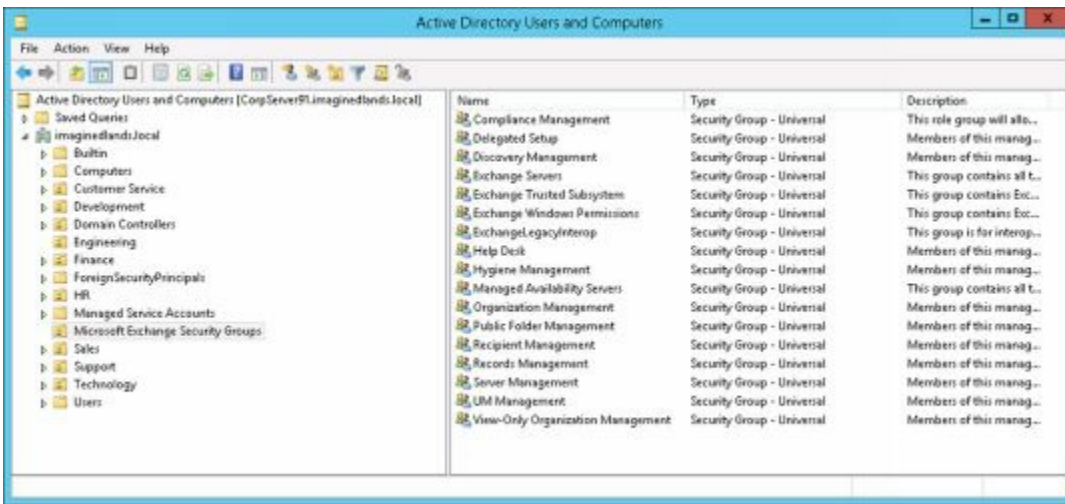


FIGURE 12-1 Using Active Directory Users And Computers to work with Exchange management groups.

In Active Directory Users And Computers, there's a hidden container of Exchange objects called Microsoft Exchange System Objects. You can display this container by selecting Advanced Features on the View menu.

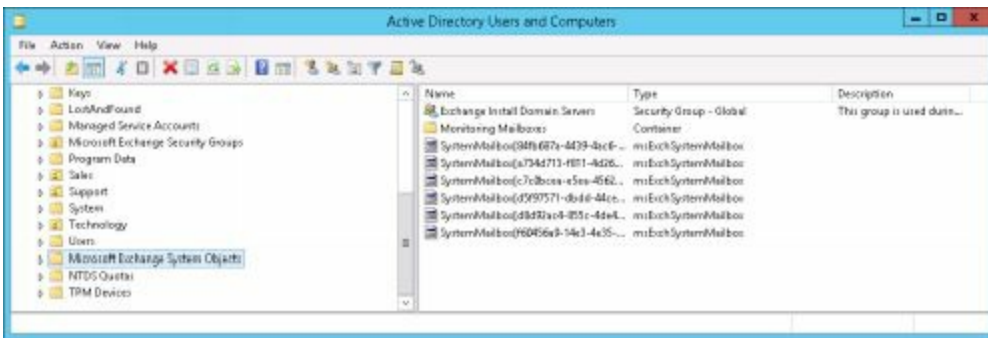


FIGURE 12-2 Viewing the Exchange system objects.

NOTE Throughout this chapter, I will often refer to Active Directory security groups simply as *security groups* or *groups*. Exchange also has distribution groups. Although distribution groups are created as objects in Active Directory, they aren't used to control access to resources.

When you are working with Exchange Online, you can view the Exchange Management groups as well. To do this, connect to Windows Azure and Microsoft Online Services in Windows PowerShell and then enter the Get-Group command. For more information on using Windows PowerShell to work with the online service, see Chapter 2 "Working with Exchange Online."

Understanding Exchange Management Groups

Table 12-1 lists predefined groups created in Active Directory for Exchange Server 2016. As the table shows, each group has a slightly different usage and purpose. Several of the groups are used by Exchange servers. These groups are Exchange Servers, Exchange Trusted Subsystem, Exchange Windows Permissions, and ExchangeLegacyInterop. You use the other groups for role-based access control and assigning management permissions. Role groups marked with an asterisk (*) are also

available with Exchange Online.

TABLE 12-1 Security groups created for Exchange 2016

GROUP	DESCRIPTION
Compliance Management*	A role group. Members of this universal security group have permission to manage compliance settings.
Delegated Setup	A role group. Members of this universal security group have permission to install and uninstall Exchange on provisioned servers.
Discovery Management*	A role group. Members of this universal security group can perform mailbox searches for data that meets specific criteria.
Exchange Install Domain Servers	Members of this global security group include domain controllers on which Exchange Server is installed. You can see this group only when you select View and then click Advanced Features in Active Directory Users And Computers.
Exchange Servers	Members of this universal security group are Exchange servers in the organization. This group allows Exchange servers to work together. By default, all computers running Exchange Server 2016 are members of this group; you should not change this setup.
Exchange Trusted Subsystem	Members of this universal security group are Exchange servers that run Exchange cmdlets using Windows Remote Management (WinRM). Members of this group have permission to read and modify all Exchange configuration settings as well as user accounts and groups.

Exchange Windows Permissions	Members of this universal security group are Exchange servers that run Exchange cmdlets using WinRM. Members of this group have permission to read and modify user accounts
------------------------------	---

	and groups.
ExchangeLegacyInterop	This group is universal security used for interoperability with Exchange Server 2003 bridgehead servers. (Shouldn't be deleted even though it is not used with Exchange 2016.)
Help Desk*	A role group. Members of this universal security group can view any property or object within the Exchange organization and have limited management permissions.
Hygiene Management*	A role group. Members of this universal security group can manage the anti-spam and antivirus features of Exchange.
Managed Availability Servers	All Mailbox servers are members of this universal security group.
Organization Management*	A role group. Members of this universal security group have full access to all Exchange properties and objects in the Exchange organization with some exceptions, such as Discovery Management.
Public Folder Management	A role group. Members of this universal security group can manage public folders and perform most public folder management operations.
Recipient Management*	A role group. Members of this universal security group have permission to modify Exchange user attributes in Active Directory and perform most mailbox operations.
Records Management*	A role group. Members of this universal security group can manage compliance features, including retention policies, message classifications, and transport rules.
Server Management	A role group. Members of this universal security group can manage all Exchange servers in the organization but do not have permission to perform global operations.

UM Management*	A role group. Members of this universal security group can manage all aspects of unified messaging (UM), including the Unified Messaging service configuration and UM recipient configuration.
View-Only Organization Management*	A role group. Members of this universal security group have read-only access to the entire Exchange organization tree in the Active Directory configuration container and read-only access to all the Windows domain containers that have Exchange recipients.

**Also available with Exchange Online*

Table 12-2 lists predefined groups and administrative roles used with Exchange Online and Office 365. These groups and roles are used for role-based access controls and assigning management permissions. However, HelpDeskAdmins and TenantAdmins aren't managed in Exchange Online. Instead, you add users to the related Office 365 role to get the desired permissions.

TABLE 12-2 Security groups and administrative roles for the Exchange Online and Office 365

GROUP/ROLE	DESCRIPTION
Billing Administrator	Used with Office 365. Members of this role are responsible for managing subscriptions and making purchases. They also can manage support tickets and monitor service health.
Exchange Administrator	Used with Office 365. Members of this role have full access to Exchange Online.
Global Administrator	Used with Office 365. Members of this role have full access to all Office 365 features and are the only ones who can assign other admin roles. Except for password admins, they also are the only ones who can reset passwords for other admins.
HelpDeskAdmins	Used with Exchange Online. Members of this group have the Password Administrator role in the Office 365 organization.
Password Administrator	Used with Office 365. Members of this role are responsible for managing passwords for standard users and other password admins. They also can manage service

	requests and monitor service health.
Service Administrator	Used with Office 365. Members of this role are responsible for managing service requests and monitoring service health.
SharePoint Administrator	Used with Office 365. Members of this role have full access to SharePoint.
Skype for Business Administrator	Used with Office 365. Members of this role can manage Skype for Business.
TenantAdmins	Used with Exchange Online. Members of this group have the Global Administrator role in the Office 365 organization.
User Management Administrator	Used with Office 365. Members of this role are responsible for managing standard users and groups. They can reset passwords for standard users, manage service requests, and monitor service health.

When working with Exchange-related groups, keep in mind that Organization Management grants the widest set of Exchange management permissions possible. Members of this group can perform any Exchange management task, including organization, server, and recipient management. Members of the Recipient Management group, on the other hand, can manage only recipient information, and Public Folder Management can manage only public folder information. View-Only Organization Management can view Exchange organization, server, and recipient information, but this group cannot manage any aspects of Exchange.

Table 12-3 provides an overview of the default group membership for the Exchange groups in an on-premises organization. Membership in a particular group grants the member the permissions of the group. Exchange groups that aren't listed don't have any default members or membership.

TABLE 12-3 Default membership for Exchange security groups

GROUP	MEMBERS	MEMBER OF
Exchange Install Domain Servers	Individual Exchange servers	Exchange Servers
Exchange Servers	Exchange Install Domain Servers, individual Exchange servers	Windows Authorization Access Group, Managed Availability Group

Exchange Trusted Subsystem	Individual Exchange servers	Exchange Windows Permissions
Exchange Windows Permissions	Exchange Trusted Subsystem	n/a
Managed Availability Servers	Exchange Servers, Mailbox servers	n/a

With Exchange Online, the TenantAdmins group is a member of the Organization Management role group and inherits its permissions from this role group. Rather than add members directly to TenantAdmins, you add members to this role by granting the Global Administrator role to users in Office 365 Admin Center.

Similarly, the HelpDeskAdmins group is a member of the View-Only Organization Management role group and inherits its permissions from this role group. Rather than add members directly to HelpDeskAdmins, you add members to this role by granting the Global Administrator role to users in Office 365 Admin Center.

Assigning Management Permissions

To grant Exchange management permissions to a user or group of users, all you need to do is make the user or group a member of the appropriate Exchange management group. For on-premises Exchange, one of the tools you can use to manage users and groups is Active Directory Users And Computers. You can make users, contacts, computers, or other group members part of an Exchange management group by completing the following steps:

1. Open Server Manager, click Tools, and then select Active Directory Users And Computers.
2. In Active Directory Users And Computers, double-click the Exchange management group you want to work with. This opens the group's Properties dialog box.
3. Click the Members tab, as shown in Figure 12-3.

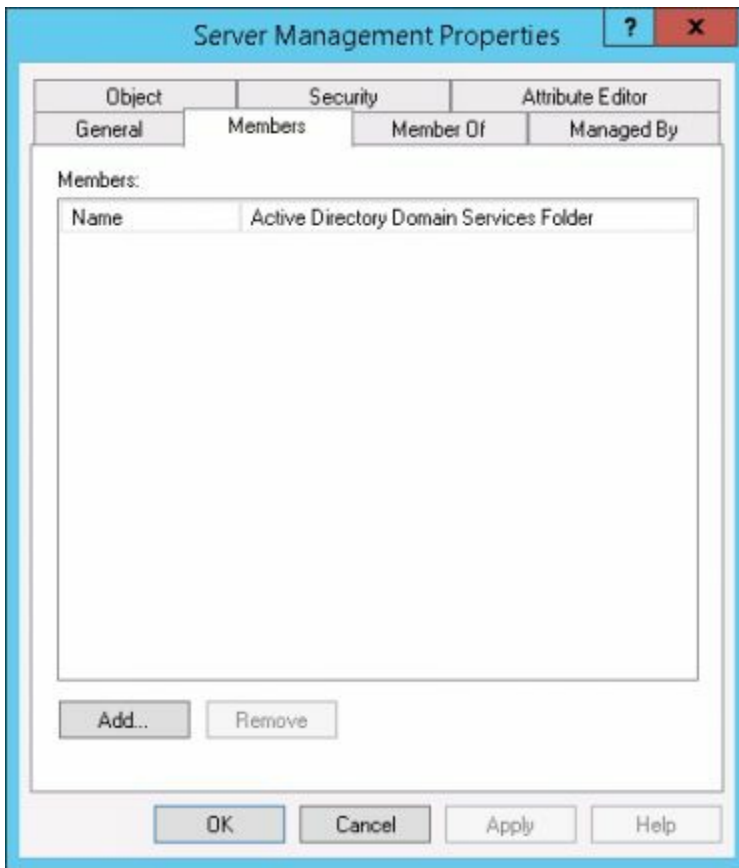


FIGURE 12-3 Using the Members tab to view and manage membership in a group.

4. To make a user or group a member of the selected group, click Add. The Select Users, Contacts, Computers, Service Accounts, Or Groups dialog box appears, as shown in Figure 12-4.


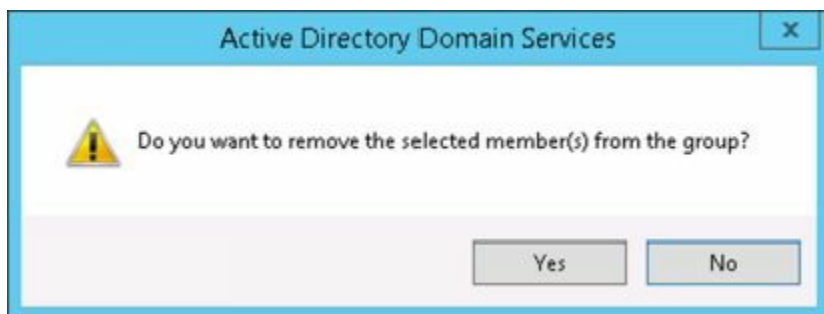


FIGURE 12-4 Specifying the name of the user, contact, computer, service account, or group to add.

5. Type the name of the account to which you want to grant permissions, and then click **Check Names**. If matches are found, select the account you want to use and then click **OK**. If no matches are found, update the name you entered, and try searching again. Repeat this step as necessary. Click **OK**.

You can remove a user, contact, computer, service account, or other group from an Exchange management group by completing the following steps:

1. In Active Directory Users And Computers, double-click the Exchange management group with you want to work with. This opens the group's Properties dialog box.
2. On the Members tab, click the user or group you want to remove and then click **Remove**. When prompted to confirm, click **Yes**, and then click **OK**.



For both on-premises Exchange and Exchange Online, you use Exchange Admin Center to manage membership in Exchange role groups. When you are managing the organization, select Permissions in the Navigation menu and then select Admin Roles to work with Exchange role groups. When you select a role, the right-most pane provides a description of the role, lists the assigned roles, and also shows the current members (see Figure 12-5). While working with this view, you can double-click a group entry to view and manage its membership.

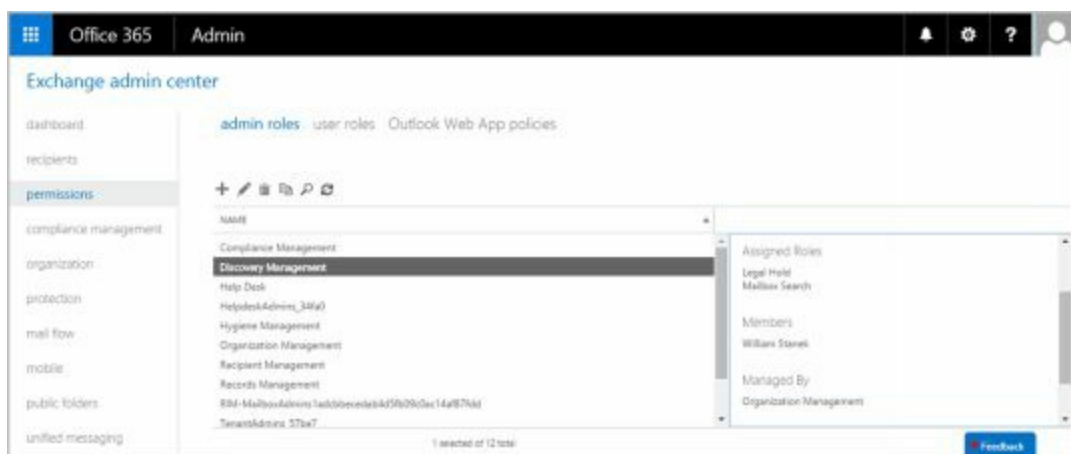
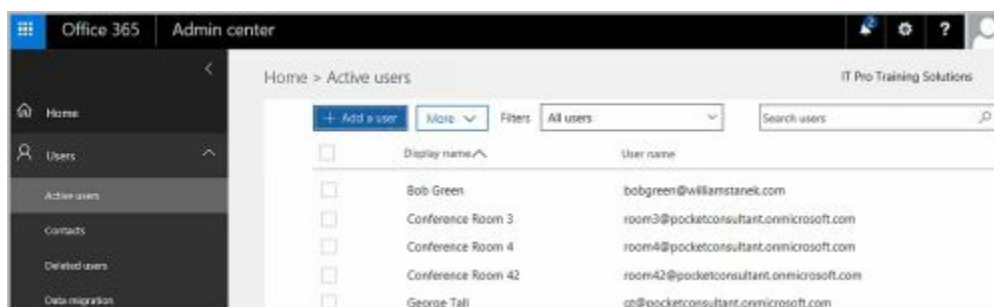


FIGURE 12-5 Using Exchange Admin Center to work with Exchange role groups.

You use Office Admin Center to manage membership in Office 365 role groups. When you are managing the Office 365 service, select **Users** in the Navigation menu and then select **Active Users** to view a list of all active users in the organization.



When you select a user, the properties page for the user is displayed. Next, select **Edit** in the Roles pane.

BO Bob Green
bobgreen@williamstane.com

Reset password Delete user

User name	bobgreen@williamstane.com	Edit
Product licenses	No products have been assigned	Edit
Group memberships (2)	Services Team Tech Resources	Edit
Sign-in status	Sign-in allowed	Edit
Roles	User (no admin access)	Edit
Display name Office phone	Bob Green	Edit

Mail Settings

More settings Manage multi-factor authentication

If you want the user to have administrator privileges, complete the following steps:

1. Choose the role to assign. For example, choose Global Administrator to make the user a member of TenantAdmins or Password Administrator to make the user a member of HelpDeskAdmins in the Exchange Online organization.
2. As necessary, enter an alternative email address for the user. Every Office 365 admin must have an alternate email address.
3. Click **Save** to apply the changes and then click **Close**.

BO Bob Green
bobgreen@williamstane.com

Edit user roles

Choose the admin role that you want to assign to this user. [Learn more about administrator roles](#)

User (no administrator access)

Global administrator

Customized administrator

- Billing administrator
- Exchange administrator
- Password administrator
- Skype for Business administrator
- Service administrator
- SharePoint administrator
- User management administrator

Alternative email address

Understanding Advanced Exchange Server Permissions

Active Directory objects are assigned a set of permissions. These permissions are standard Microsoft Windows permissions, object-specific permissions, and extended permissions.

Table 12-4 summarizes the most common object permissions. Keep in mind that some permissions are generalized. For example, with Read Value(s) and Write Value(s), Value(s) is a placeholder for the actual type of value or values.

TABLE 12-4 Common permissions for Active Directory objects

PERMISSION	DESCRIPTION
Full Control	Permits reading, writing, modifying, and deleting
List Contents	Permits viewing object contents
Read All Properties	Permits reading all properties of an object
Write All Properties	Permits writing to all properties of an object
Read Value(s)	Permits reading the specified value(s) of an object, such as general information or group membership
Write Value(s)	Permits writing the specified value(s) of an object, such as general information or group membership
Read Permissions	Permits reading object permissions
Modify Permissions	Permits modifying object permissions
Delete	Permits deleting an object
Delete Subtree	Permits deleting the object and its child objects
Modify Owner	Permits changing the ownership of the object
All Validated Writes	Permits all types of validated

	writes
All Extended Writes	Permits all extended writes
Create All Child Objects	Permits creating all child objects
Delete All Child Objects	Permits deleting all child objects
Add/Remove Self As Member	Permits adding and removing the object as a member
Send To	Permits sending to the object
Send As	Permits sending as the object
Change Password	Permits changing the password for the object
Receive As	Permits receiving as the object

Table 12-5 summarizes Exchange-specific permissions for objects. If you want to learn more about other types of permissions, I recommend that you read *Windows Server 2016: Essentials for Administration* (Stanek & Associates, 2016).

TABLE 12-5 Extended permissions for Exchange Server

PERMISSION	DESCRIPTION
Read Exchange Information	Permits reading general Exchange properties of the object
Write Exchange Information	Permits writing general Exchange properties of the object
Read Exchange Personal Information	Permits reading personal identification and contact information for an object
Write Exchange Personal Information	Permits writing personal identification and contact information for an object
Read Phone and Mail Options	Permits reading phone and mail options of an object
Write Phone and	Permits writing phone and mail

Although you can use standard Windows permissions, object-specific permissions, and extended permissions to control Exchange management and use, Microsoft recommends that you use role-based access controls instead. My recommendation is to use the role-based access controls whenever possible in place of specific permissions.

Assigning Advanced Exchange Server Permissions

In Active Directory, different types of objects can have different sets of permissions. Different objects can also have general permissions that are specific to the container in which they're defined. For troubleshooting or fine-tuning your environment, you might occasionally need to modify advanced permissions. You can set advanced permissions for Active Directory objects by following these steps:

1. Open Active Directory Users And Computers. If advanced features aren't currently being displayed, select **Advanced Features** on the View menu.
2. Right-click the user, group, service account, or computer account with which you want to work.

CAUTION Only administrators with a solid understanding of Active Directory and Active Directory permissions should manipulate advanced object permissions. Incorrectly setting advanced object permissions can cause problems that are difficult to track down and may also cause irreparable harm to the Exchange organization.

3. Select **Properties** from the shortcut menu, and then click the Security tab in the Properties dialog box, as shown in Figure 12-6.

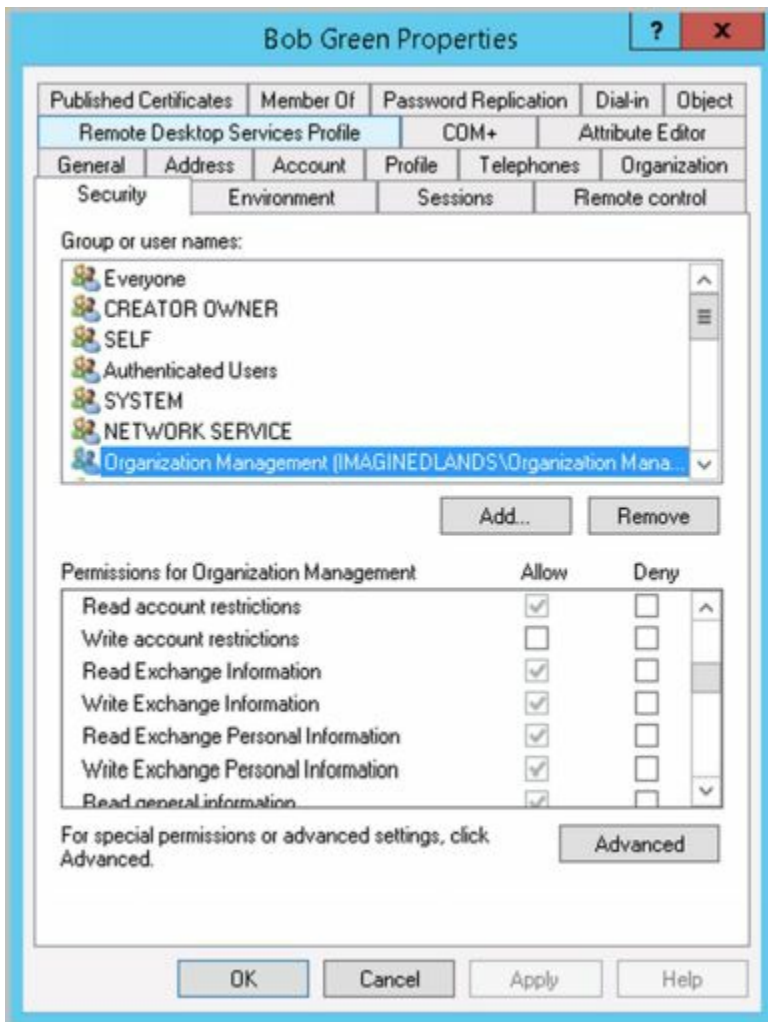


FIGURE 12-6 Using the Security tab to manage advanced permissions.

4. Users or groups with access permissions are listed in the Group Or User Names list box. You can change permissions for these users and groups by doing the following:
 - Select the user or group you want to change.
 - Use the Permissions list box to grant or deny access permissions.
 - When inherited permissions are dimmed, override inherited permissions by selecting the opposite permissions.
5. To set access permissions for additional users, computers, or groups, click **Add**. Then use the Select Users, Computers, Security Accounts, Or Groups dialog box to add users, computers, security accounts, or groups.
6. Select the user, computer, service account, or group you want to configure in the Group Or User Names list box, click **Add**, and then click **OK**. Then use the fields in the Permissions area to allow or deny permissions. Repeat this step for other users, computers, service accounts, or groups. Click **OK** when you're finished.

Configuring Role-Based Permissions for Exchange

Exchange 2016 and Exchange Online implement role-based access controls that allow you to easily customize permissions for users in the organization. You use role-based access controls to do the following:

- [Assign permissions to groups of users](#)
- [Define policies that assign permissions](#)
- [Assign permissions directly to users](#)

Before I discuss each of these tasks, I'll discuss essential concepts related to role-based permissions. Because the permissions model is fairly complex, I recommend reading this entire section to understand your implementation options before starting to assign permissions.

Understanding Role-Based Permissions

Role-based access control is a permissions model that uses role assignment to define the management tasks a user or group of users can perform in the Exchange organization. Exchange defines many built-in management roles that you can use to manage your Exchange organization. Each built-in role acts as a logical grouping of permissions that specify the management actions that those assigned the role can perform. You also can create custom roles.

You can assign roles to role groups or directly to users. You also can assign roles through role policies that are then applied to role groups, users, or both. By assigning roles, you grant permission to perform management tasks.

At the top of the permissions model is the role group, which is a special type of security group that has been assigned one or more roles. Keep the following in mind when working with role-based permissions:

- [You can assign role-based permissions to any mailbox-enabled user account.](#)
[Assigning a role to a user grants the user the ability to perform a specific management action.](#)
- [You can assign role-based permissions to any universal security group.](#) [Assigning a role to a group grants members of the group the ability to perform a specific management action.](#)
- [You cannot assign role-based permissions to security groups with the domain local or global scope.](#)
- [You cannot assign role-based permissions to distribution groups regardless of scope.](#)

As Table 12-1 showed previously, Exchange 2016 and Exchange Online include a number of predefined role groups. These role groups are assigned fixed management roles by default. As a result, you do not need to explicitly add roles to these groups to enable management, nor can you add or remove roles associated with the built-in groups. You can, however, manage the members of the predefined role groups using the

procedures discussed previously. You can also create your own role groups and manage the membership of those groups.

When you assign a role to a group, the management scope determines where in the Active Directory hierarchy that objects can be managed by users assigned a management role. The scope is either implicitly or explicitly assigned. Implicit scopes are the default scopes that apply based on a particular type of management role.

Table 12-6 lists key management roles with an organization scope. A role with an organization scope applies across the whole Exchange organization. Table 12-7 lists key management roles with an organization scope that apply to individual servers. Table 12-8 lists key management roles with a user scope. A role with a user scope applies to an individual user. When you create a role group, you also can set an explicit scope, such as for objects in the Customer Service organizational unit or objects in the Technology organizational unit.

TABLE 12-6 Management roles with an organization scope

MANAGEMENT ROLE	ENABLES MANAGERS TO...
Active Directory Permissions	Configure Active Directory permissions in an organization. Keep in mind that permissions set directly on Active Directory objects cannot be enforced through RBAC.
Address Lists	Manage address lists, the global address list, and offline address lists in an organization.
Audit Logs	Manage audit logs in an organization.
Cmdlet Extension Agents	Manage cmdlet extension agents in an organization.
Data Loss Prevention	Configure data loss prevention settings in an organization.
Database Availability Groups	Manage database availability groups in an organization.
Disaster Recovery	Restore mailboxes and database availability groups in an organization.
Distribution Groups	Create and manage distribution groups and distribution group members in an organization.
Edge Subscriptions	Manage edge synchronization and subscription configuration between Edge Transport servers and Mailbox servers in an organization.

E-Mail Address Policies	Manage email address policies in an organization.
Exchange Connectors	Manage routing group connectors, delivery agent connectors, and other connectors used for transport. This role doesn't enable administrators to manage Send and Receive connectors.
Federated Sharing	Manage cross-forest and cross-organization sharing in an organization.
Information Rights Management	Manage the Information Rights Management (IRM) features of Exchange in an organization.
Journaling	Manage journaling configuration in an organization.
Legal Hold	Configure whether data within a mailbox should be retained for litigation purposes in an organization.
Mail Enabled Public Folders	Configure whether individual public folders are mail enabled or mail disabled in an organization.
Mail Recipient Creation	Create mailboxes, mail users, mail contacts, distribution groups, and dynamic distribution groups in an organization.
Mail Recipients	Manage existing mailboxes, mail users, mail contacts, distribution groups, and dynamic distribution groups in an organization. This does not enable administrators to create these recipients.
Mail Tips	Manage mail tips in an organization.
Message Tracking	Track messages in an organization.
Monitoring	Monitor the Microsoft Exchange services and component availability in an organization.
Move Mailboxes	Move mailboxes between servers in an organization and between servers in the local organization and another organization.
Organization Configuration	Manage basic organization-wide settings. This role type doesn't include the permissions included in the Organization Client Access or Organization Transport Settings role types.
Organization Transport	

Settings	Manage organization-wide transport settings, including system messages, site configuration, and so forth. This role doesn't enable administrators to create or manage transport Receive or Send connectors, queues, hygiene, agents, remote and accepted domains, or rules.
Public Folders	Manage public folders in an organization. This role type doesn't enable administrators to manage whether public folders are mail enabled or to manage public folder replication.
Send Connectors	Manage transport send connectors in an organization.
Recipient Policies	Manage recipient policies, such as provisioning policies, in an organization.
Remote and Accepted Domains	Manage remote and accepted domains in an organization.
Reset Password	Reset users' password in an organization.
Retention Management	Manage retention policies in an organization.
Role Management	Manage management role groups, role assignment policies, management roles, role entries, assignments, and scopes in an organization. Users assigned roles associated with this role type can override the Managed By property for role groups, configure any role group, and add or remove members to or from any role group.
Security Group Creation and Membership	Create and manage security groups and their memberships in an organization.
Team Mailboxes	Define site mailbox provisioning policies and manage site mailboxes.
Transport Agents	Manage transport agents in an organization.
Transport Hygiene	Manage antivirus and anti-spam features in an organization.
Transport Rules	Manage transport rules.
UM Mailboxes	Manage the unified messaging (UM) configuration of mailboxes and other recipients.
UM Prompts	Create and manage custom UM voice prompts.

Unified Messaging	Manage Unified Messaging settings. This role doesn't enable administrators to manage UM-specific mailbox configuration or UM prompts.
User Options	View the Microsoft Outlook Web Access options for users.
View-Only Configuration	View all of the nonrecipient Exchange configuration settings.
View-Only Recipients	View the configuration of recipients, including mailboxes, mail users, mail contacts, distribution groups, and dynamic distribution groups.
View-Only Audit Logs	Search the administrator audit logs and view results.

TABLE 12-7 Management roles for individual servers

MANAGEMENT ROLE	ENABLES MANAGERS TO...
Database Copies	Manage mailbox database copies on individual servers.
Databases	Create, manage, mount, and dismount mailbox and public folder databases on individual servers.
Exchange Server Certificates	Create, import, export, and manage Exchange server certificates on individual servers.
Exchange Servers	Manage Exchange server configuration on individual servers.
Exchange Virtual Directories	Manage Autodiscover, Outlook Web App, Exchange ActiveSync, offline address book (OAB), Windows PowerShell, and Web administration interface virtual directories on individual servers.
Migration	Migrate mailboxes and mailbox content into or out of a server.
POP3 and IMAP4 Protocols	Manage Post Office Protocol version 3 (POP3) and Internet Message Access Protocol version 4 (IMAP4) configuration, such as authentication and connection settings, on individual servers.

Receive Connectors	Manage transport Receive connector configuration, such as size limits on an individual server.
Transport Queues	Manage transport queues on an individual server.

TABLE 12-8 Management roles for user scope

MANAGEMENT ROLE	ENABLES INDIVIDUAL USERS TO...
MyBaseOptions	View and modify the basic configuration of their own mailboxes and associated settings.
MyContactInformation	Modify their contact information. This information includes their addresses and phone numbers.
MyDiagnostics	Perform basic diagnostics on their mailboxes.
MyDistributionGroupMembership	View and modify their membership in distribution groups in an organization, provided that those distribution groups allow manipulation of group membership.
MyDistributionGroups	Create, modify, and view distribution groups and modify, view, remove, and add members to distribution groups they own.
MyProfileInformation	Modify their names.
MyRetentionPolicies	View their retention tags and view and modify their retention tag settings and defaults.
MyTeamMailboxes	Create and connect site mailboxes.
MyTextMessaging	View and modify their text messaging settings.
MyVoiceMail	View and modify their voice mail settings.

Role assignment policies grant users permissions to configure their Outlook Web App options and perform limited management tasks. When you install Exchange 2016, the setup process creates the Default Role Assignment Policy and sets this as the default for all new mailboxes. This policy grants users the MyBaseOptions, MyContactInformation, MyDistributionGroupMembership, and MyVoiceMail roles, but it does not grant users the MyDistributionGroups and MyProfileInformation roles.

Exchange Online has a Default Role Assignment policy as well. This default policy,

assigned to all Exchange Online users, grants all of the management roles. You can create other role assignment policies as well.

Working with Role Groups

By default, members of the Organization Management group can manage any role group in the Exchange organization. Anyone designated as a manager of a role group can manage the role group. You assign a user as a manager of a role group using the `-ManagedBy` parameter, which can be set when you create or modify a role group.

To view the currently available role groups and the roles they've been assigned, select **Permissions** in the Navigation menu and then select **Admin Roles**. As shown in Figure 12-7, when you select a role group, the details pane lists the assigned roles and members.

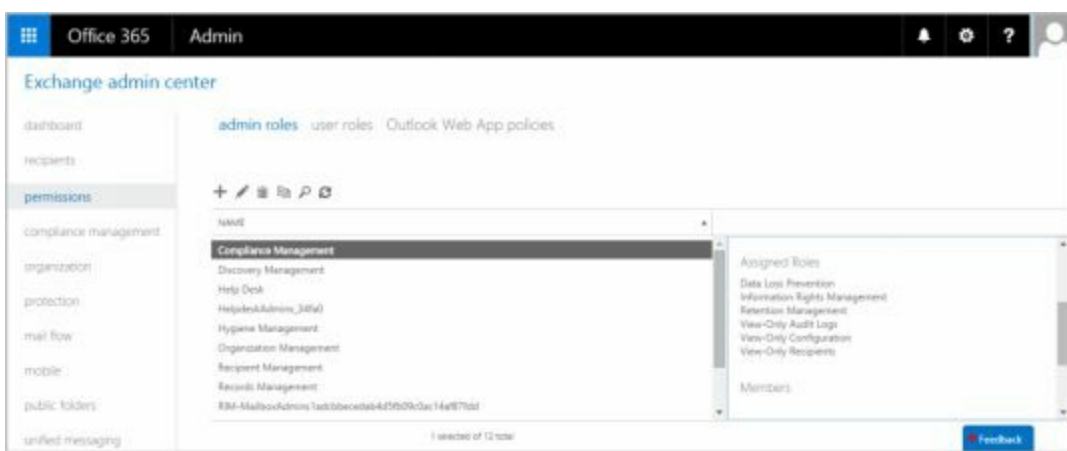



FIGURE 12-7 Viewing the role groups and the assigned roles and members of a selected group.

To create a role group, complete the following steps:

1. In Exchange Admin Center, select **Permissions** in the Navigation menu and then select **Admin Roles**.
2. Click **New** (). In the New Role Group dialog box, shown in Figure 12-8, type a descriptive name for the role group. By default, the role group will use the implicit write scope.

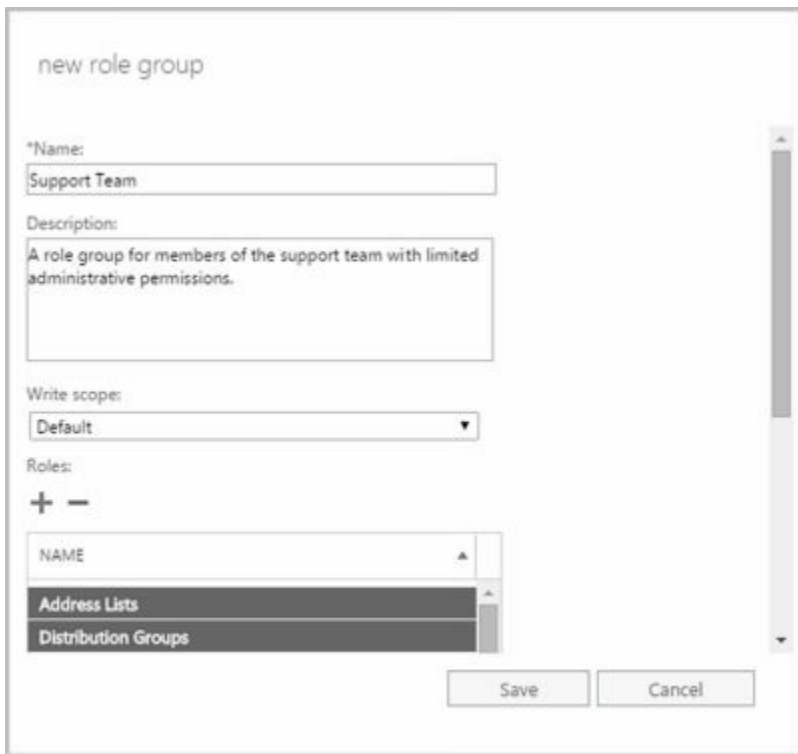

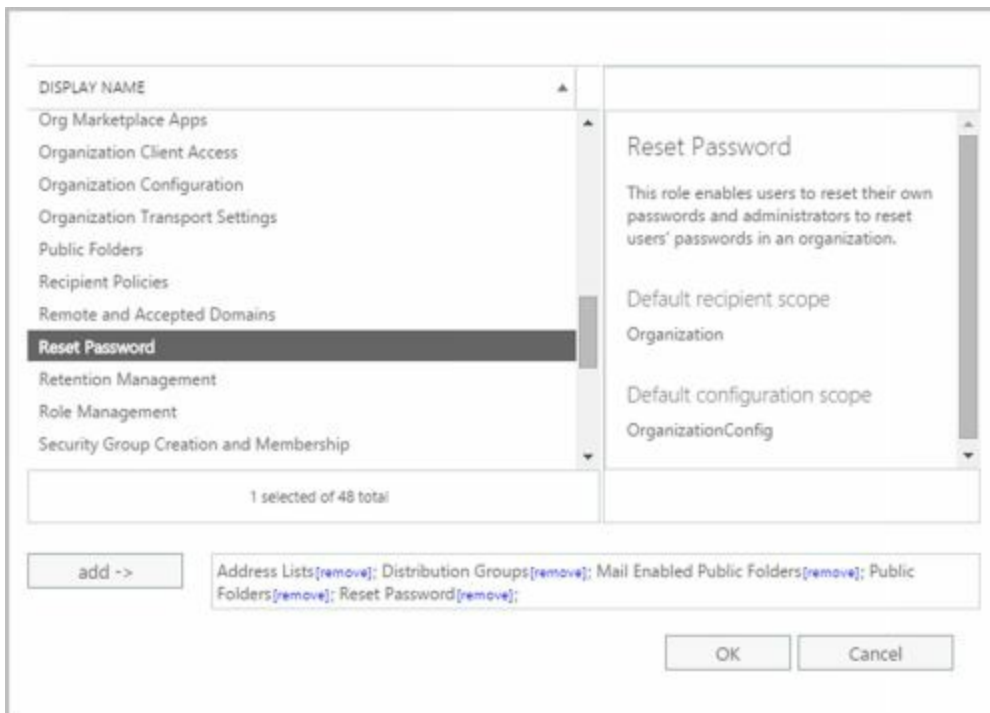



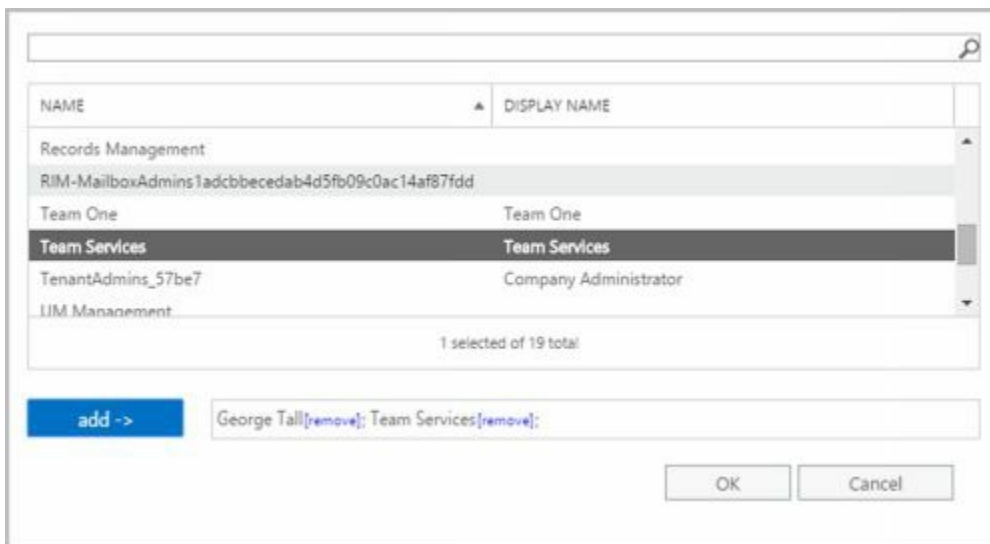
FIGURE 12-8 Creating a new role group.

- Under Roles, click Add (). In the Select A Role dialog box select roles to assign to the role group and then click **Add**. You can select multiple roles using the Shift or Ctrl key, or you can simply select and add each role individually. When you are finished adding roles, click OK.



- Under Members, click Add (). In the Select Members dialog box select members to add to the role group and then click **Add**. You can select multiple members using the Shift or Ctrl key, or you can simply select and add each member individually. When you are finished, click OK.

5. Click **Save** to create the role group.



In the shell, commands you use to work with role groups include the following:

- **Get-RoleGroup** Displays a complete or filtered list of role groups. When specifying filters, use parentheses to define the filter, such as **-Filter { RolegroupType -Eq " Linked " }**.

```
Get-RoleGroup [-Identity RoleGroupName ] {AddtlParams}
```

```
{AddtlParams}
```

```
[-AccountPartition PartitionID ] [-DomainController FullyQualifiedName ]  
[-Filter {LinkedGroup | ManagedBy | Members | Name | RoleGroupType |  
DisplayName}] [-Organization OrganizationID ] [-ReadFromDomainController  
{$True|$False}] [-ResultSize Size ] [-SortBy {LinkedGroup |  
ManagedBy | Members | Name | RoleGroupType | DisplayName}]  
[-ShowPartnerLinked {$True|$False}] [-UsnForReconciliationSearch Num ]
```

- **New-RoleGroup** Creates a new role group. When specifying roles, you must use the full role name, including spaces. Enclose the role names in quotation marks and separate each role with a comma, such as “ **Mail Recipient Creation** ”, “ **Mail Recipients** ”, “ **Recipient Policies** ”.

```
New-RoleGroup -Name RoleGroupName [-Roles Roles ]  
[-ManagedBy ManagerIds ] [-Members MemberIds ] {AddtlParams}
```

```
{AddtlParams}
```

```
[-CustomConfigWriteScope Scope ] [-CustomRecipientWriteScope Scope ]  
[-Description Description ] [-DisplayName DisplayName ]  
[-DomainController FQDN ] [-ExternalDirectoryObjectId ObjId ]  
[-Organization OrganizationID ] [-PartnerManaged {$True|$False}]  
[-RecipientOrganizationalUnitScope Scope ]  
[-SamAccountName PreWin2000Name ] [-ValidationOrganization OrgId ]  
[-WellKnownObjectGUID GUID ]
```

```
[-LinkedCredential Credential ] [-LinkedDomainController LinkedDC ]
```

`[-LinkedForeignGroup LinkedGroup]`

- **Remove-RoleGroup** Removes a role group. If a role group has designated managers, you must be listed as a manager to remove the role group or use the `-BypassSecurityGroupManagerCheck` parameter and be an organization manager.

`Remove-RoleGroup -Identity RoleGroupName {AddtlParams}`

`{AddtlParams}`

`[-BypassSecurityGroupManagerCheck {$True|$False}]`

`[-DomainController FullyQualifiedName] [-ForReconciliation`

`{$True|$False}] [-RemoveWellKnownObjectGUID {$True|$False}]`

- **Set-RoleGroup** Configures role group properties. If you specify managers, you must provide the complete list of managers because the list you provide overwrites the existing list of managers. To manage role assignment, see the “Assigning Roles Directly or Via Policy” section later in the chapter.

`Set-RoleGroup -Identity RoleGroupName [-ManagedBy ManagerIds]`

`[-Name NewName] {AddtlParams}`

`{AddtlParams}`

`[-BypassSecurityGroupManagerCheck {$True|$False}]`

`[-Description Description] [-DisplayName DisplayName]`

`[-DomainController FullyQualifiedName]`

`[-ExternalDirectoryObjectId ObjId]`

`[-LinkedCredential Credential] [-LinkedDomainController LinkedDC]`

`[-LinkedForeignGroup LinkedGroup]`

You use `New-RoleGroup` to create role groups. When you create a role group, you must specify the group name and the roles assigned to the group. You should also specify the managers and members of the group. The managers and members can be individual users or groups identified by their display name, alias, or distinguished name. If you want to specify more than one manager or member, separate each entry with a comma. The following example creates the Special Recipient Management role group to allow members of the group to manage (but not create) recipients:

```
New-RoleGroup -Name "Special Recipient Management"
```

```
-Roles "mail recipients", "recipient policies"
```

```
-ManagedBy "juliec", "tylerk", "ulij"
```

```
-Member "mikeg", "lylep", "rubyc", "yus"
```

By default, the scope of the role group is the organization. You can also set a specific scope for an organizational unit. The following example creates a role group named LA Recipient Management and sets the scope to the LA Office organizational unit to allow members of the group to manage recipients in the LA Office organizational unit:

```
New-RoleGroup -Name "LA Recipient Management"
```

```
-Roles "mail recipient creation", "mail recipients", "recipient policies"
```


-ManagedBy "LA Managers" -Member "LA Help Desk"
-RecipientOrganizationalUnitScope "LA Office"

A linked role group links the role group to a universal security group in another forest. Creating a linked role group is useful if your Exchange servers reside in a resource forest and your users and managers reside in a separate user forest. If you create a linked role group, you can't add members directly to it. You must add the members to the universal security group in the foreign forest.

When you create linked role groups, you use the -LinkedDomainController parameter to specify the fully qualified domain name or IP address of a domain controller in the foreign forest. This domain controller is used to get security information for the foreign universal security group, which is specified by the -LinkedForeignGroup parameter. If you use the -LinkedDomainController parameter, you must specify a foreign universal security group with the -LinkedForeignGroup parameter, and you can't use the -Members parameter. Optionally, you can use the -LinkedCredential parameter to specify credentials to use to access the foreign forest. To pass in the credentials, use a Credential object.

The following example creates a linked role group that enables the members of the Chicago Managers universal security group to manage recipients located in the Chicago office:

```
$cred = Get-Credentials
```

```
New-RoleGroup -Name "Chicago Recipient Managers"  
-LinkedDomainController corpserver26.cpusers.imagedlands.com  
-LinkedCredential $cred -LinkedForeignGroup "Chicago Managers"  
-CustomRecipientWriteScope "Chicago Recipients" -Roles "mail recipients"
```

In this example, Chicago Managers is a group created in the user forest and the administrator is logged on to the resource forest. When PowerShell reads the Get-Credentials command, a prompt for the user name and password for the user forest appears.

Role groups are created as universal security groups in the Active Directory database. In Active Directory Users And Computers, you'll find role groups in the Microsoft Exchange Security Groups container. After you create a role group, you can manage it using Active Directory Users And Computers or Exchange Management Shell. The management tasks you can perform depend on which tool you are using. In Active Directory Users And Computers, you can manage group membership, rename the group, or delete the group. Additional tasks you can perform when you use Exchange Management Shell include setting managers and modifying role assignments.

NOTE Although you can edit a group's managers or other attributes in Active Directory Users And Computers, you shouldn't do this because some values are linked and set differently than you'd expect. For example, you set the ManagedBy property to the distinguished name of the first manager and define additional

managers using the `msExchCoManagedByLink` property.

You can list available role groups using `Get-RoleGroup`. If you type `Get-RoleGroup` at the Exchange Management Shell prompt, you see a list of all role groups defined in the Exchange organization to which you are connected. You can filter the output in a variety of ways using standard PowerShell filtering techniques. `Get-RoleGroup` also has a `-Filter` parameter that you can use to filter the output according to specific criteria you set. The following example looks for a role group named CS Recipient Management and lists all its properties:

```
Get-RoleGroup -filter {Name -eq "CS Recipient Management"} |  
format-list
```

You can use `Set-RoleGroup` to change the name of a role group or to define a new list of managers. To delete a role group, use `Remove-RoleGroup`.

Managing Role Group Members

By default, members of the Organization Management group can manage the membership of any role group in the Exchange organization. Anyone designated as a manager of a role group can manage the membership of that role group as well.

In the shell, commands you use to configure role group membership include the following:

- **Add-RoleGroupMember** Adds a user or universal security group as a member of a role group. If a role group has designated managers, you must be listed as a manager to add role group members or use the `-BypassSecurityGroupManagerCheck` parameter and be an organization manager.

```
Add-RoleGroupMember -Identity RoleGroupName -Member MemberIds  
[-BypassSecurityGroupManagerCheck {$True|$False}]  
[-DomainController FullyQualifiedName ]
```

- **Get-RoleGroupMember** Lists the members of a role group.

```
Get-RoleGroupMember -Identity RoleGroupName  
[-DomainController FullyQualifiedName ]  
[-ReadFromDomainController {$True|$False}]  
[-ResultSize Size ]
```

- **Remove-RoleGroupMember** Removes a user or universal security group from a role group. If a role group has designated managers, you must be listed as a manager to remove role group members or use the `-BypassSecurityGroupManagerCheck` parameter and be an organization manager.

```
Remove-RoleGroupMember -Identity RoleGroupName -Member MemberIds  
[-BypassSecurityGroupManagerCheck {$True|$False}]  
[-DomainController FullyQualifiedName ]
```

- **Update-RoleGroupMember** Replaces the current group membership with the list of members you provide.

```
Update-RoleGroupMember -Identity RoleGroupName -Members NewMemberIds  
[-BypassSecurityGroupManagerCheck {$True|$False}]  
[-DomainController FullyQualifiedName ] @techjob
```

You add members to a role group using `Add-RoleGroupMember`. When you add a member to a role group, the member is given the effective permissions provided by the management roles assigned to the role group. If the role group has designated managers, you must be a role group manager or use the `-BypassSecurityGroupManagerCheck` parameter to override the security group management check. The following example adds a user to the LA Recipient Management role group:

```
Add-RoleGroupMember -Identity "LA Recipient Management"  
-Member "joym"
```

Whether you are working with Exchange Online or on-premises Exchange at the shell prompt, don't forget that all the features of PowerShell are at your disposal. The following example lists all users with mailboxes in the Technology department and adds them to the Technology Management role group:

```
Get-User -Filter { Department -Eq "Technology" -And -RecipientType  
-Eq "UserMailbox" } | Get-Mailbox | Add-RoleGroupMember  
"Technology Management"
```

You can list members of a particular role group using `Get-RoleGroupMember`. Members are listed by name and recipient type as shown in the following example and sample output:

```
Get-RoleGroupMember -Identity "CS Recipient Management"
```

Name	RecipientType
-----	-----
Riis Anders	UserMailbox
Darren Waite	UserMailbox

You can delete role group members using `Remove-RoleGroupMember`. When you remove a member from a role group, the user or group of users can no longer perform the management tasks made available by that role group. However, keep in mind that the user or group of users might be a member of another role group that grants management permissions. If so, the user or group of users will still be able to perform management tasks.

NOTE For linked role groups, you can't use `Remove-RoleGroupMember` to remove members from the role group. Instead, you need to remove members from the foreign universal security group (USG) that's linked to the linked role group. Use `Get-RoleGroup` to identify the foreign group.

Assigning Roles Directly or Via Policy

You can assign built-in or custom roles to users, role groups, and universal security groups in one of two ways:

- [Directly using role assignment](#)
- [Via assignment policy](#)

Directly assigning roles is accomplished using role assignment commands. By adding, removing, or modifying role assignments, you can control the management tasks that users can perform. Although you can assign roles directly to users or universal security groups, this approach increases the complexity of the permissions model in your Exchange organization. A more flexible solution is to assign roles via assignment policy. Assigning roles via assignment policy requires you to do the following:

1. [Create assignment policies.](#)
2. [Assign roles to these policies.](#)
3. [Assign policies to users or groups as appropriate.](#)

Management roles define the specific tasks that can be performed by the members of a role group assigned the role. A role assignment links a management role and a role group. Assigning a management role to a role group grants members of the role group the ability to perform the management tasks defined in the management role. Role assignments can use management scopes to control where the assignment can be used.

In the shell, commands you use to work with role assignment include the following:

- [Get-ManagementRoleAssignment](#) Displays a complete or filtered list of role assignments for a role group. You can examine role assignments by name, assignment type, or scope type as well as whether the assignment is enabled or disabled.

```
Get-ManagementRoleAssignment [-Identity RoleAssignmentToRetrieve ]
{AddtlParams}
```

```
Get-ManagementRoleAssignment [-Role RoleID ] [-RoleAssignee IdentityToCheck ]
[-AssignmentMethod {Direct | SecurityGroup |
RoleAssignmentPolicy}] {AddtlParams}
```

```
{AddtlParams}
[-ConfigWriteScope <None | NotApplicable | OrganizationConfig |
CustomConfigScope | PartnerDelegatedTenantScope |
ExclusiveConfigScope>] [-CustomConfigWriteScope ManagementScopeId ]
[-CustomRecipientWriteScope ManagementScopeId ] [-Delegating <$true
| $false>] [-DomainController FullyQualifiedName ] [-Enabled <$true
| $false>] [-Exclusive <$true | $false>]
[-ExclusiveConfigWriteScope ManagementScopeId ]
[-ExclusiveRecipientWriteScope ManagementScopeId ]
[-GetEffectiveUsers <$true | $false>]
[-GetEffectiveUsers <$true | $false>]
[-Organization OrganizationId ] [-RecipientOrganizationalUnitScope
OrganizationalUnitId ] [-RecipientWriteScope <None | NotApplicable
| Organization | MyGAL | Self | MyDirectReports | OU |
CustomRecipientScope | MyDistributionGroups | MyExecutive |
ExclusiveRecipientScope>] [-RoleAssigneeType <User |
```

SecurityGroup | RoleAssignmentPolicy | MailboxPlan | ForeignSecurityPrincipal | RoleGroup | LinkedRoleGroup>] [-WritableDatabase **DatabaseId**] [-WritableRecipient **GeneralRecipientId**] [-WritableServer **ServerId**]

- **New-ManagementRoleAssignment** Creates a new role assignment, and assigns it directly to a user or group or assigns it via an assignment policy.

New-ManagementRoleAssignment -Name **RoleAssignmentName** -SecurityGroup **Group** -Role **Roles** {AddtlParams}

New-ManagementRoleAssignment -Name **RoleAssignmentName** -Policy **Policy** -Role **Roles** {AddtlParams}

New-ManagementRoleAssignment -Name **RoleAssignmentName** -User **User** -Role **Roles** {AddtlParams}

New-ManagementRoleAssignment -Name **RoleAssignmentName** -Computer **Computer** -Role **Roles** {AddtlParams}

{AddtlParams}
[-CustomConfigWriteScope **Scope**] [-CustomRecipientWriteScope **Scope**] [-Delegating {\$True|\$False}] [-DomainController **FullyQualifiedName**] [-ExclusiveConfigWriteScope **Scope**] [-ExclusiveRecipientWriteScope **Scope**] [-Organization **OrganizationId**] [-RecipientOrganizationalUnitScope **Scope**] [-RecipientRelativeWriteScope <None | NotApplicable | Organization | MyGAL | Self | MyDirectReports | OU | CustomRecipientScope | MyDistributionGroups | MyExecutive | ExclusiveRecipientScope>] [-UnscopedTopLevel {\$True|\$False}]

- **Remove-ManagementRoleAssignment** Removes a role assignment.

Remove-ManagementRoleAssignment -Identity **RoleAssignmentName** [-DomainController **FullyQualifiedName**]

- **Set-ManagementRoleAssignment** Configures role assignment properties.

Set-ManagementRoleAssignment -Identity **RoleAssignmentName** [-DomainController **FullyQualifiedName**] [-Enabled {\$True|\$False}] {AddtlParams1 | AddtlParams2 | AddtlParams3 | AddtlParams4}

{AddtlParams1}
[-CustomConfigWriteScope **Scope**] [-RecipientOrganizationalUnitScope **OUIId**] [-RecipientRelativeWriteScope <None | NotApplicable | Organization | MyGAL | Self | MyDirectReports | OU | CustomRecipientScope | MyDistributionGroups | MyExecutive | ExclusiveRecipientScope>]

{AddtlParams2}
[-CustomConfigWriteScope **Scope**]

`[-CustomRecipientWriteScope Scope]`

`{AddtlParams3}`

`[-CustomConfigWriteScope Scope]`

`[-DomainController FullyQualifiedName]`

`{AddtlParams4}`

`[-ExclusiveConfigWriteScope Scope]`

`[-ExclusiveRecipientWriteScope Scope]`

You can list role assignments using `Get-ManagementRoleAssignment`. You use `New-ManagementRoleAssignment` to assign roles. The following example assigns the Retention Management role to the Central Help Desk group:

```
New-ManagementRoleAssignment -Name "Central Help Desk_Retention"  
-Role "Retention Management" -SecurityGroup "Central Help Desk"
```

The following example assigns the Mail Recipients role to members of the Marketing Help Desk group and restricts the write scope to the Marketing organizational unit:

```
New-ManagementRoleAssignment -Name "Marketing_Options"  
-Role "Mail Recipients" -SecurityGroup "Marketing Help Desk"  
-RecipientOrganizationalUnitScope "imaginedlands.com/Marketing"
```

This allows users who are members of the Marketing Help Desk group to manage existing mailboxes, mail users, mail contacts, distribution groups, and dynamic distribution groups in the Marketing organizational unit. This does not enable these users to create recipients in this organizational unit. To create recipients, the users need to be assigned the Mail Recipient Creation role.

You can modify role assignment using `Set-ManagementRoleAssignment`. The following example disables the Central Help Desk_Retention role assignment:

```
Set-ManagementRoleAssignment -Identity "Central Help Desk_Retention"  
-Enabled $False
```

When you disable a role assignment, the users assigned the role can no longer perform the management tasks granted by the role. However, keep in mind that a user might have been granted the permission in another way. By disabling a role assignment rather than removing it, you can easily enable the role assignment again as shown in the following example:

```
Set-ManagementRoleAssignment -Identity "Central Help Desk_Retention"  
-Enabled $True
```

However, if you are sure you no longer want to use a particular role assignment, you can remove it using `Remove-ManagementRoleAssignment` as shown in the following example:

```
Remove-ManagementRoleAssignment -Identity "Central Help Desk_Retention"
```

When you create a new assignment policy, you can assign it to users using the `New-`

Mailbox, Set-Mailbox, or Enable-Mailbox cmdlet. If you make the new assignment policy the default assignment policy, it's assigned to all new mailboxes that don't have an explicitly designated assignment policy. After you create an assignment policy, you must assign it at least one management role for permissions to apply to a mailbox. Without any roles assigned to it, users assigned the policy won't be able to manage any of their mailbox configurations. To assign a management role, use New-ManagementRoleAssignment.

In the shell, commands you use to work with role assignment policy include the following:

- **Get-RoleAssignmentPolicy** Lists all policies or a specified role assignment policy.

```
Get-RoleAssignmentPolicy [-Identity AssignmentPolicyName ]  
[-DomainController FullyQualifiedName ] [-Organization OrganizationId ]
```

- **New-RoleAssignmentPolicy** Creates a new role assignment policy.

```
New-RoleAssignmentPolicy -Name AssignmentPolicyName  
[-Description Description ] [-DomainController FullyQualifiedName ]  
[-IsDefault {$True|$False}] [-Organization OrganizationId ]
```

- **Remove-RoleAssignmentPolicy** Removes a role assignment policy.

```
Remove-RoleAssignmentPolicy -Identity AssignmentPolicyName  
[-DomainController FullyQualifiedName ]
```

- **Set-RoleAssignmentPolicy** Changes the name of a role assignment policy, or sets a role assignment policy as the default.

```
Set-RoleAssignmentPolicy -Identity AssignmentPolicyName  
[-Description Description ] [-DomainController FullyQualifiedName ]  
[-IsDefault {$True|$False}] [-Name NewName ]
```

You can list role assignment policies using Get-RoleAssignmentPolicy. Rather than view all available assignment policies, you can easily filter the output to look for default assignment policies. Here is an example:

```
Get-RoleAssignmentPolicy | Where { $_.IsDefault -eq $True }
```

You use New-RoleAssignmentPolicy to create role assignment policies. The following example creates the Standard User Policy and assigns it as the default:

```
New-RoleAssignmentPolicy -Name "Standard User Policy"
```

When you create a new assignment policy, you can assign it to users using New-Mailbox, Set-Mailbox, or Enable-Mailbox as shown in the following example:

```
Set-Mailbox -Identity "tommyj" -RoleAssignmentPolicy "Standard User  
Policy"
```

If you make the new assignment policy the default assignment policy, it's assigned to all

new mailboxes that don't have an explicitly designated assignment policy. You can specify that a policy is the default when you create it using `-IsDefault`. You can also designate a policy as the default using `Set-RoleAssignmentPolicy` as shown in this example:

```
Set-RoleAssignmentPolicy -Identity "Standard User Policy" -IsDefault
```

After you create an assignment policy, you must assign at least one management role to it for it to apply permissions to a mailbox. Without any roles assigned to it, users assigned the policy won't be able to manage any of their mailbox configuration. To assign a management role, use `New-ManagementRoleAssignment`.

You can remove policies using `Remove-RoleAssignmentPolicy`. The assignment policy you want to remove can't be assigned to any mailboxes or management roles. Also, if you want to remove the default assignment policy, it must be the last assignment policy. Because of this, you need to use `Set-Mailbox` to change the assignment policy for any mailbox that's assigned the assignment policy before you can remove it. If the assignment policy is the default assignment policy, use `Set-RoleAssignmentPolicy` to select a new default assignment policy before you remove the old default policy. You don't need to do this if you're removing the last assignment policy. Additionally, keep in mind that you can use `Remove-ManagementRoleAssignment` to remove any management role assignments assigned to a policy.

With this in mind, the following series of examples show how you can modify and remove assignment policy. The first example removes the assignment policy called "Standard User Policy" by finding all of the mailboxes assigned the policy and then assigning a different policy:

```
Get-Mailbox | Where {$_.RoleAssignmentPolicy -Eq "Standard User Policy"}  
| Set-Mailbox -RoleAssignmentPolicy "New User Policy"
```

Next, you can remove all the role assignments assigned to an assignment policy:

```
Get-ManagementRoleAssignment -RoleAssignee "Standard User Policy" |  
Remove-ManagementRoleAssignment
```

Afterward, you can remove the assignment policy by entering the following:

```
Remove-RoleAssignmentPolicy "Standard User Policy"
```

Configuring Account Management Permissions

Exchange 2016 and Exchange Online user roles control the settings that users can configure on their own mailboxes and on distribution groups they own. These settings determine whether users can:


- [Change the display name, contact information, text messaging settings, voice mail settings, and more.](#)
- [View and modify apps, mail subscriptions, and retention policies.](#)
- [Modify the basic configuration of the mailbox.](#)
- [Create and connect site mailboxes.](#)

- Manage text messaging and voice mail settings
- Create, modify, and view distribution groups
- Manage membership of distribution groups they own.
- Manage their membership in distribution groups.

The Exchange organization has a default role assignment policy that grants users permission to configure all user-manageable settings. You can create one or more additional role assignment policies and assign them to users at any time using Exchange Admin Center. To view the currently available policies, select **Permissions** in the Navigation menu and then select **User Roles** as shown in Figure 12-9.



FIGURE 12-9 Configuring user roles to manage permissions.

To create a policy, click **New** (). In the Role Assignment Policy dialog box, type a descriptive name for the policy, such as All Standard Users. To grant a role to users, select the related check box. To not grant a role to users, clear the related check box. At a minimum, be sure to grant MyBaseOptions to the policy so that those assigned the policy can access their mailbox and basic settings.

role assignment policy

*Name:

Description:

Contact information:

MyContactInformation
 This role enables individual users to modify their contact information, including address and phone numbers.

MyAddressInformation

MyMobileInformation

MyPersonalInformation

Finally, click **Save** to create the policy and update the organization settings. It may take several minutes to update the organization settings. If an error occurs, try to create the policy again before you begin any troubleshooting. Sometimes, a complex process won't be completed fully the first time and retrying will resolve the problem.

To assign a policy to a user, follow these steps:

1. In Exchange Admin Center, select **Recipients** in the Navigation menu and then select **Mailboxes**. Double-click the entry for the user.
2. On the Mailbox Features page, use the Role Assignment Policy selection list to choose the policy that you want to apply.
3. Click **Save**.

Jeff Peterson

general

mailbox usage

contact information

organization

email address

▶ mailbox features

member of

MailTip

mailbox delegation

Select the mailbox settings, phone and voice features, and email connectivity options for this mailbox. [Learn more](#)

Sharing policy:

Role assignment policy:

Retention policy:

Address book policy:

Managing Advanced Permissions

Advanced permissions areas you can work with are related to custom management roles, management scopes, and role entries. Management roles define the management tasks users can perform. Management scopes identify the objects that are allowed to be managed. Role entries are the individual permission entries on a management role that allow users to perform management tasks.

Adding Custom Roles

The built-in roles were listed previously in Tables 12-6 to 12-8. The built-in roles are fixed, and you cannot create role entries to define additional management tasks for built-in roles. You can, however, create your own custom roles based on built-in roles and then extend the custom roles as necessary to meet the needs of your organization. In this way, custom management roles allow you to do things you can't do with the built-in roles.

Commands you use to create custom roles and to view any existing roles include the following:

- **Get-ManagementRole** Displays a complete or filtered list of management roles defined in the organization. Role types are the same as those listed previously without spaces in their names.

```
Get-ManagementRole [-Identity RoleName] [-DomainController  
FullyQualifiedName] [-Organization OrganizationId] [-RoleType RoleType]  
{AddtlParams}
```

```
{AddtlParams}  
{ [-Cmdlet Cmdlet] [-CmdletParameters Parameters] |  
[-GetChildren {$True|$False}] |  
[-Script Script] [-ScriptParameters Parameters] |  
[-Recurse {$True|$False}] }
```

- **New-ManagementRole** Creates a new management role.

```
New-ManagementRole -Name RoleName  
[-Parent ParentRoleToCopy | -UnScopedTopLevel {$True|$False}]  
[-Description Description] [-DomainController FullyQualifiedName]  
[-Organization OrganizationId]
```

- **Remove-ManagementRole** Removes a management role.

```
Remove-ManagementRole [-Identity RoleName]  
[-DomainController FullyQualifiedName] [-Recurse {$True|$False}]  
[-UnScopedTopLevel {$True|$False}]
```

To view management roles, you use `Get-ManagementRole`. Entering `Get-ManagementRole` by itself without parameters lists all the roles in your organization.

Additional options include using

- **-Identity** to view information about a specific role
- **-Cmdlet** to list all roles that include a specified cmdlet
- **-CmdletParameters** to list all roles that include the specified cmdlet parameter or parameters
- **-GetChildren** to list only the child roles of a specified parent role
- **-Recurse** to list the role specified in the **-Identity** parameter, its child roles, and all subsequent children until all the roles that were created based on the parent role have been fully identified.
- **-RoleType** to list all roles of a particular type
- **-Script** to list all roles that include a specified script
- **-ScriptParameters** to list all roles that include the specified script parameter or parameters

The following example lists all the roles associated with the Mail Recipient Creation role:

```
Get-ManagementRole "Mail Recipient Creation" -Recurse
```

You can create your own custom roles using **New-ManagementRole**. New roles can either be empty top-level roles or based on an existing parent role. For example, the following command creates an empty role:

```
New-ManagementRole -Name "Change Management"  
-UnscopedTopLevel
```

In the following example, a new role is created based on the Organization Client Access role:

```
New-ManagementRole -Name "Organization Client Access View-Only"  
-Parent "Organization Client Access"
```

After you create a role based on another role, you might need to remove role entries that are not required. For example, the following command ensures the Organization Client Access View-Only role grants view-only permission for Client Access information by removing any entries for commands that don't begin with **Get**:

```
Get-ManagementRoleEntry "Organization Client Access View-Only\*" |  
Where { $_.Name -NotLike "Get*" } | Remove-ManagementRoleEntry
```

To remove a custom role, you use **Remove-ManagementRole**. You can remove a role by name as shown in the following example:

```
Remove-ManagementRole "Organization Client Access View-Only"
```

Using the **-Recurse** parameter, you can remove all child roles of a role. Using the **-UnscopedTopLevel** parameter, you can remove an unscoped top-level role. You also can use **Get-ManagementRole** to obtain a list of roles to remove as shown in this example:

```
Get-ManagementRole *MyTestRole* | Remove-ManagementRole
```

TIP To avoid accidentally removing a number of important roles, you should run `Get-ManagementRole` by itself first or add the `-WhatIf` parameter to `Remove-ManagementRole`. Either technique will ensure you know exactly which roles you are working with.

Adding Custom Role Scopes

Every management role has a management scope that determines where in Active Directory objects can be viewed or modified by users assigned the management role. Management scopes can be defined as either regular or exclusive. Regular scopes can be either implicitly or explicitly created. They are simply the standard type of scope, and they define the set of recipients that can be managed. Exclusive scopes, on the other hand, must always be explicitly created, and they allow you to deny users access to objects contained within the exclusive scope if those users aren't assigned a role associated with the exclusive scope.

Scopes can be:

- Inherited from the management role
- Specified as a predefined relative scope on a management role assignment
- Created using custom filters and added to a management role assignment

Scopes inherited from management roles are called *implicit scopes*, while predefined and custom scopes are called *explicit scopes*. Implicit scopes include:

- **Recipient read scope** Determines which recipient objects the user assigned the management role is allowed to read from Active Directory.
- **Recipient write scope** Determines which recipient objects the user assigned the management role is allowed to modify in Active Directory.
- **Configuration read scope** Determines which configuration objects the user assigned the management role is allowed to read from Active Directory.
- **Configuration write scope** Determines which organizational and server objects the user assigned the management role is allowed to modify in Active Directory.

Commands you use to work with scopes include the following:

- **Get-ManagementScope** Displays a complete or filtered list of management scopes defined in the organization.

```
Get-ManagementScope [-Identity ScopeName ]  
[-Exclusive {$True|$False}] [-DomainController FullyQualifiedName ]  
[-Organization OrganizationId ] [-Orphan {$True|$False}]
```

- **New-ManagementScope** Creates a new management scope.

```
New-ManagementScope -Name ScopeName -RecipientRestrictionFilter  
Filter [-RecipientRoot Root ] {AddtlParams}
```

```
New-ManagementScope -Name ScopeName  
-ServerList Servers | -ServerRestrictionFilter Filter {AddtlParams}
```

New-ManagementScope -Name **ScopeName**
-DatabaseList **Servers** | -DatabaseRestrictionFilter **Filter** {AddtlParams}

{AddtlParams}
[-DomainController **FullyQualifiedName**] [-Organization **OrganizationId**]
[-Exclusive {\$True|\$False}] [-Force {\$True|\$False}]

- **Remove-ManagementScope** Removes a management scope.

Remove-ManagementScope [-Identity **Scope**]
[-DomainController **FullyQualifiedName**]

- **Set-ManagementScope** Modifies the settings of a management scope.

Set-ManagementScope -Identity **ScopeName** -ServerRestrictionFilter
Filter [-DomainController **FullyQualifiedName**] [-Name **Name**]

Set-ManagementScope -Identity **ScopeName** -RecipientRestrictionFilter
Filter [-RecipientRoot **Root**] [-DomainController **FullyQualifiedName**]
[-Name **Name**]

Set-ManagementScope -Identity **ScopeName** -DatabaseRestrictionFilter
Filter [-DomainController **FullyQualifiedName**] [-Name **Name**]

You use Get-ManagementScope to retrieve a list of existing management scopes. If you want to list only exclusive scopes, use the -Exclusive parameter. If you want to list only management scopes that aren't associated with role assignments, use the -Orphan parameter, as shown here:

Get-ManagementScope -Orphan

You can create custom management scopes using New-ManagementScope. After you create a regular or exclusive scope, you need to associate the scope with a management role assignment. One way to do this is to use New-ManagementRoleAssignment.

You define scopes using recipient restriction filters, explicit server lists, or server restriction filters. For example, the following command creates the Sales Team scope that applies only to mailboxes located in the Sales organizational unit:

```
New-ManagementScope -Name "Sales Team Scope" -RecipientRoot  
"imaginedlands.com/Sales" -RecipientRestrictionFilter {RecipientType -eq  
"UserMailbox"}
```

The following example creates a scope that applies only to MailServer14 and MailServer22:

```
New-ManagementScope -Name "Main Server Scope" -ServerList  
"MailServer14", "MailServer22"
```

The following example creates a scope that applies only to servers in the Active Directory site called Seattle-First-Site:

```
New-ManagementScope -Name "Seattle Site Scope" -ServerRestrictionFilter  
{ServerSite -eq "Seattle-First-Site"}
```

Exclusive scopes work a bit differently. When an exclusive scope is created, all users are immediately blocked from modifying the recipients that match the exclusive scope until the scope is associated with a management role assignment. If other role assignments are associated with other exclusive scopes that match the same recipients, those assignments can still modify the recipients. For example, the following command creates a Protected Managers exclusive scope for users that contain the string “Manager” in their job titles:

```
New-ManagementScope -Name "Protected Managers"  
-RecipientRestrictionFilter { Title -Like "*Manager*" } -Exclusive
```

After creating an exclusive scope, you then need to associate it with a management role assignment that assigns the appropriate management roles to the appropriate role group or groups. In the following example, members of the Level 5 Administrators security group are granted permission to work with Protected Manager mailboxes:

```
New-ManagementRoleAssignment -Name "Level 5 Administrators_Mail  
Recipients" -SecurityGroup "Level 5 Administrators" -Role "Mail  
Recipients" -CustomRecipientWriteScope "Protected Managers"
```

You use Set-ManagementScope to modify the settings of a management scope. If you change a scope that has been associated with management role assignments, the updated scope applies to all of the associated role assignments. To remove a management scope, you can use Remove-ManagementScope. However, you can't remove a management scope if it's associated with a role assignment.

Adding Custom Role Entries

Role entries determine the management actions that members of a role group can perform. You create a role entry by specifying the permitted management command and any permitted command parameters.

Assigning a management role to a role group is essentially similar to creating the related role entries that allow a user or group to perform related management tasks. Another way to grant permission to perform a management action is to create a management role entry and add it to a management role. However, keep in mind that you can't add role entries to built-in roles.

Commands you use to work with role entries include:

- **Add-ManagementRoleEntry** Adds role entries to a custom management role. You can't add role entries to built-in roles. The **-UnScopedTopLevel** parameter allows you to specify that you're adding a custom script or non-Exchange cmdlet to an unscoped top-level management role.

```
Add-ManagementRoleEntry -Identity RoleEntryToAdd  
[-DomainController FullyQualifiedName ] [-Parameters CmdletParametersToUse ]
```



```
[-PSSnapinName SnapinThatContainsCmdlet ] [-Type <Cmdlet | Script |  
ApplicationPermission | All>] [-Overwrite {$True|$False}]  
[-UnScopedTopLevel {$True|$False}]
```

```
Add-ManagementRoleEntry -ParentRoleEntry ParentRoleEntry  
-Role Role [-DomainController FullyQualifiedName ]  
[-Overwrite {$True|$False}]
```

- **Get-ManagementRoleEntry** Lists the role entries configured on a particular role. You can list role entries that match specific criteria such as role name, cmdlet name, parameter name, role entry type, or associated PowerShell snap-in.

```
Get-ManagementRoleEntry -Identity RoleEntry  
[-DomainController FullyQualifiedName ]  
[-Parameters CmdletParameters ] [-PSSnapinName Snapin ]  
[-Type <Cmdlet | Script | ApplicationPermission | All>]
```

- **Remove-ManagementRoleEntry** Removes a management role entry.

```
Remove-ManagementRoleEntry -Identity RoleEntry  
[-DomainController FullyQualifiedName ]
```

- **Set-ManagementRoleEntry** Modifies a management role entry.

```
Set-ManagementRoleEntry -Identity RoleEntry  
[-AddParameter {$True|$False} | -RemoveParameter {$True|$False}]  
[-Parameters ParametersToAddOrRemove ]
```

```
[-DomainController FullyQualifiedName ]  
[-UnScopedTopLevel {$True|$False}]
```

Every management role must have at least one management role entry. A role entry consists of a single cmdlet and its parameters, a script, or a special permission that you want to make available. If a cmdlet or script doesn't appear as an entry on a management role, that cmdlet or script isn't accessible via that role. Similarly, if a parameter isn't specified in a role entry, the parameter on that cmdlet or script isn't accessible via that role.

The way you create and work with role entries depends on whether they are based on the built-in roles or unscoped roles. Roles based on built-in roles can contain only role entries that are Exchange cmdlets. To use custom scripts or non-Exchange cmdlets, you need to add them as unscoped role entries to an unscoped top-level role.

You can't add management role entries to child roles if the entries don't appear in parent roles. For example, if the parent role doesn't have an entry for New-Mailbox, the child role can't be assigned that cmdlet. Additionally, if Set-Mailbox is on the parent role but the -Database parameter has been removed from the entry, the -Database parameter on the Set-Mailbox cmdlet can't be added to the entry on the child role. With this in mind, you need to carefully choose the parent role to copy when you want to create a new customized role.

Role entry names are a combination of the management role that they're associated with and the name of the cmdlet or script that you want to make available. The role name and the cmdlet or script are separated by a backslash character (\). For example, the role entry name for the New-Mailbox cmdlet on the Mail Recipient Creation role is Mail Recipient Creation\New-Mailbox.

You can use the wildcard character (*) in the role entry name to return all of the role entries that match the input you provide. The wildcard character can be used with role names as well as with cmdlet or script names. For example, you can use * to return a list of all role entries for all roles, *\New-Mailbox to return a list of all role entries that contain the New-Mailbox cmdlet, or Mail Recipient Creation* to return a list of all role entries on the Mail Recipient Creation role.

When you create a role entry, you need to specify all of the parameters that can be used. Exchange will try to verify the parameters that you provide when you add the role entry. Only the parameters that you include are available to the users assigned to the role. You need to update role entries manually if parameters available for cmdlets or scripts change.

To avoid errors, keep the following in mind:

- [Scripts that you add to an unscoped role entry must reside in the Exchange 2016 scripts directory on every server where administrators and users connect using Exchange Management Shell. The default scripts directory is C:\Program Files\Microsoft\Exchange Server\V15\Scripts.](#)
- [Non-Exchange cmdlets that you add to an unscoped role entry must be installed on every Exchange 2016 server where administrators and users connect using the Exchange Management Shell. When you add a non-Exchange cmdlet, you must specify the Windows PowerShell snap-in name that contains the non-Exchange cmdlet.](#)

You use Get-ManagementRoleEntry to list role entries that have been configured on roles. For example, the following command lists all the role entries that exist on the Mail Recipient Creation role:

```
Get-ManagementRoleEntry "Mail Recipient Creation\*"
```

You also can list all the role entries that contain a particular command, as shown here:

```
Get-ManagementRoleEntry *\Get-Recipient
```

You can list role entries that match specific criteria such as role name or cmdlet name. Using Add-ManagementRoleEntry, you can specify role entries to add to a role. You specify the role entry to add using the -Identity parameter and the basic syntax for the identity as RoleName\CmdletName. Role entries are either based on a parent role entry or are unscoped (the default), specified using the -ParentRoleEntry or -UnScopedTopLevel parameter, respectively. The -Role parameter specifies the role to which the new role entry is added.

For example, the following command adds a role entry for the Get-Mailbox cmdlet to

the LA Recipient Managers role:

```
Add-ManagementRoleEntry -Identity "LA Recipient Managers\Get-Mailbox"
```

This entry assigns permission for the Get-Mailbox cmdlet to members of the LA Recipient Managers role. You can specify the exact parameters that are permitted as shown in the following example:

```
Add-ManagementRoleEntry -Identity "LA Recipient Managers\Get-Mailbox"  
-Parameters Archive, Identity, Filter, OrganizationalUnit, SortBy
```

You can also assign permission for multiple commands. Consider the following example:

```
Get-ManagementRoleEntry "Mail Recipients\Get-Mailbox*" |  
Add-ManagementRoleEntry -Role "Central Help Desk"
```

Here, Get-ManagementRoleEntry is used to retrieve a list of all the role entries for the Mail Recipients role that begin with the string “Get-Mailbox” in the cmdlet name, and then add them to the Central Help Desk role using the Add-ManagementRoleEntry cmdlet. The role entries are added to the child role exactly as they’re configured on the parent role, Mail Recipients.

You use Set-ManagementRoleEntry to change the available parameters on an existing management role entry. With the -AddParameter parameter, the parameters you specify are added to the role entry. With the -RemoveParameter parameter, the parameters you specify are removed from the role entry. Otherwise, only the parameters you specify are included in the role entry. For example, with Get-Mailbox you might want users to be able to specify a server and limit the result set size, and you can do this by adding the -Server and -ResultSize parameters as shown in this example:

```
Set-ManagementRoleEntry -Identity "LA Recipient Managers\Get-Mailbox"  
-AddParameter Server, ResultSize
```

To remove all parameters, set -Parameters to \$Null and don’t use either -AddParameter or -RemoveParameter as shown in this example:

```
Set-ManagementRoleEntry -Identity "LA Recipient Managers\Get-Mailbox"  
-Parameters $Null
```

You use Remove-ManagementRoleEntry to remove role entries. However, you can’t remove role entries from built-in management roles.

Working with Shared and Split Permissions

When you deploy Exchange 2016, you can use a shared permissions model or one of two split permissions models. Which permissions model your organization uses depends squarely on who should have the right to create and manage security principals in Active Directory.

Using Shared Permissions

The shared permissions model is the default. With the shared permissions model, management of Exchange and Active Directory are not separated within the Exchange management tools. Administrators can use the Exchange management tools to create security principals in Active Directory. In this model, the Mail Recipient Creation role allows administrators to create security principals, such as Active Directory users, and the Security Group Creation And Membership role allows administrators to create security groups and manage security group membership.

Two Exchange role groups have these roles by default:

- The Organization Management role group has the Mail Recipient Creation role and the Security Group Creation And Membership role. This means members of this role group can create users, security groups, and other security principals in Active Directory. They also can manage security group membership.
- The Recipient Management role group has the Mail Recipient Creation role. This means members of this role group can create security principals in Active Directory, but cannot create security groups or manage the membership of security groups.

If you want other users to be able to create security principals and manage the membership of security groups, you have several choices. You can assign the Mail Recipient Creation role, the Security Group Creation And Membership role, or both roles to other role groups, users, and security groups. You also can make the appropriate users, security groups, or both members of the appropriate role group.

IMPORTANT Permissions for working with security groups are separated from permissions for working with other security principals because Exchange administrators typically don't need to be able to create or manage security groups. In fact, in the base model, anyone who needs to be able to create or manage security groups is assumed to be an advanced administrator or manager who requires organization-wide management permissions.

An option for extending the shared permissions model is to grant the Security Group Creation And Membership role to the Recipient Management role group. This approach:

- Allows members of the Recipient Management role group to create and manage security groups in Active Directory.
- Doesn't require granting the role to individual users and security groups as may be

needed for management of the Exchange organization.

I recommend this configuration only when Exchange administrators need to create security groups as part of their regular routine. With this option, you can continue to grant the Mail Recipient Creation role, the Security Group Creation And Membership role, or both roles to other role groups, users, and security groups as well.

Using Split Permissions

Some organizations require strict management of who can create security principals, and this is where split permissions are useful. With split permissions, you remove the default settings that allow members of Recipient Management and Organization Management to create security principals in Active Directory. Thereafter the process of creating security principals and the process of configuring Exchange attributes for security principals are completely separate. As a result, Active Directory administrators are responsible for creating security principals and Exchange administrators are responsible for configuring the Exchange attributes associated with security principals.

With split permissions, you have two configuration options. You can use:

- **RBAC split permissions** With RBAC split permissions, only those who are members of the appropriate role groups can create Active Directory security principals and manage group membership.
- **Active Directory split permissions** With Active Directory split permissions, permissions to create and manage security principals and group membership are not available in the Exchange management tools. You must use Active Directory management tools to create and manage security principals.

TIP For organizations that require split permissions, Microsoft recommends using RBAC split permissions and so do I. With RBAC split permissions, you can continue to use the Exchange management tools to create and manage security principals in Active Directory, and this gives you more flexibility in how you can use and work with Exchange.

Each Exchange organization has one and only one permissions model. Your Exchange organization is either configured to use a shared model that allows for RBAC split permissions or it's configured to use Active Directory split permissions. During installation of Exchange 2016, you can specify whether you want to use Active Directory split permissions. If you select this option, the shared permissions and RBAC split permissions models are not available.

To move between the shared model that allows for RBAC split permissions and the Active Directory split permissions model or vice versa, you must run the following command from the Exchange 2016 installation media:

```
setup.exe /PrepareAD /ActiveDirectorySplitPermissions: {$true|$false}
```

where \$true sets the organization to use Active Directory split permissions and \$false sets the organization to use the shared model that allows for RBAC split permissions.

You have to prepare Active Directory in each instance because many changes to groups and group membership will be made in the background. Next, you must either wait for Active Directory to replicate an access token to all servers running Exchange 2010 or Exchange 2016, or you must restart all servers running Exchange 2010 or Exchange 2016. Finally, you must implement your permissions model. A step-by-step procedure with examples follows:

1. Create a role group for Active Directory administrators and assign the Mail Recipient Creation role and the Security Group Creation And Membership role to this role group. If you want members of this role group to be able to create role assignments, include the Role Management role. Complete this step by adding members to the new role group.

```
New-RoleGroup "AD Admins" -Roles "Mail Recipient Creation",  
"Security Group Creation and Membership", "Role Management"  
Add-RoleGroupMember "AD Admins" -Member williams, timb, anneh, mikel
```

2. If you want members of the new role group to be able to delegate any of the roles they've been assigned, you can create delegating assignments.

```
New-ManagementRoleAssignment -Role "Mail Recipient Creation"  
-SecurityGroup "AD Admins" -Delegating
```

```
New-ManagementRoleAssignment -Role "Security Group Creation and  
Membership" -SecurityGroup "AD Admins" -Delegating
```

3. If you only want members of the new role group to be able to manage the group membership, replace the delegate list on the role group.

```
Set-RoleGroup "Active Directory Administrators" -ManagedBy  
"AD Admins"
```

4. If you are implementing RBAC split permissions, remove the Mail Recipient Creation role and the Security Group Creation And Membership role assignments from the Recipient Management and Organization Management role groups.

```
Get-ManagementRoleAssignment -Role "Mail Recipient Creation" | Where  
{ $_.RoleAssigneeName -eq "Recipient Management" or  
$_.RoleAssigneeName -eq "Organization Management"} |  
Remove-ManagementRoleAssignment -Whatif
```

```
Get-ManagementRoleAssignment -Role " Security Group Creation and  
Membership " | Where { $_.RoleAssigneeName -eq "Recipient Management"  
or $_.RoleAssigneeName -eq "Organization Management"} |  
Remove-ManagementRoleAssignment -Whatif
```

CAUTION I recommend running the commands in the step with the `-Whatif` parameter first. This will ensure the command does exactly what you think it will. Before you remove these roles, confirm that the new role group has been assigned these roles and that the new role group has the appropriate members. Your account should be a member of the new role group.

5. Determine what groups have been assigned the Mail Recipient Creation role and the Security Group Creation And Membership role. Optionally, remove the Mail Recipient Creation role and the Security Group Creation And Membership role assignments from all other users and groups.

```
Get-ManagementRoleAssignment -Role *Creation* | Format-List Name, Role,  
RoleAssigneeName
```

```
Get-ManagementRoleAssignment -Role "Mail Recipient Creation" | Where  
{ $_.RoleAssigneeName -NE "AD Admins" } |  
Remove-ManagementRoleAssignment -Whatif
```

```
Get-ManagementRoleAssignment -Role "Security Group Creation and  
Membership" | Where { $_.RoleAssigneeName -NE "AD Admins" } |  
Remove-ManagementRoleAssignment -Whatif
```

When you use split permissions, only members of the group created in the previous procedure will be able to use the Exchange management tools to:

- Create mailbox users, mail-enabled users, mail-enabled contacts, remote mailbox users, and security groups.
- Remove mailbox users, mail-enabled users, mail-enabled contacts, remote mailbox users, and security groups.

This means Exchange administrators and others won't be able to use `New-Mailbox`, `New-MailContact`, `New-MailUser`, `New-RemoteMailbox`, `Remove-Mailbox`, `Remove-MailContact`, `Remove-MailUser`, or `Remove-RemoteMailbox`. Additionally, with Active Directory split permissions, only members of the group will be able to create distribution groups and manage their membership. Thus, only members of the group will be able to use the following cmdlets:

- `Add-DistributionGroupMember`, `New-DistributionGroup`
- `Remove-DistributionGroup`, `Remove-DistributionGroupMember`
- `Update-DistributionGroupMember`

Exchange administrators will still be able to configure Exchange attributes on existing Active Directory security principals. They will also be able to create and manage Exchange-specific objects.

Chapter 13. Implementing Exchange Services

You can implement Exchange services in several ways, including:

- **On-premises** With an on-premises implementation, you deploy Exchange server hardware on your network and manage all aspects of the implementation, including server configuration, organization configuration, and recipient configuration.
- **Online** With an online (or cloud-only) implementation, you rely on hardware and services provided by Microsoft. All aspects of the server configuration are managed by Microsoft. You manage the service-level settings, organization configuration, and recipient configuration.
- **Hybrid** With a hybrid implementation, you integrate on-premises and online implementations. The on-premises and Exchange Online organizations use a shared domain namespace, so mail is securely routed between them, and you can easily share data between the implementations.

When you use an online implementation, Microsoft manages the hardware configuration and ensures availability. Otherwise, you are responsible for any on-premises hardware.

In terms of functionality, Exchange Server 2016 is an incremental release, building on the radical changes in Exchange Server 2010 and adding enhancements to the refinements found in Exchange Server 2013. Like Exchange Server 2010 and Exchange Server 2013, Exchange Server 2016 does away with the concepts of storage groups, Local Continuous Replication (LCR), Single Copy Clusters (SCC), and clustered mailbox servers. This means that:

- Databases are no longer associated with storage groups.
- Database availability groups are used to group databases for high availability.
- Databases are managed at the organization level instead of at the server level.

Exchange Server 2016 integrates high availability into the core architecture by enhancing aspects of Cluster Continuous Replication (CCR) and Standby Continuous Replication (SCR) and combining them into a single, high-availability solution for both on-site and off-site data replication.

Exchange Server 2016 also provides for automatic failover and recovery without requiring clusters when you deploy multiple mailbox servers. Because of these changes, building a high-availability mailbox server solution doesn't require cluster hardware or advanced cluster configuration. Instead, database availability groups provide the base component for high availability. Failover is automatic for mailbox databases that are part of the same database availability group.

The basic rules for database availability groups have not changed since implementation in Exchange Server 2010. Each mailbox server can have multiple databases, and each database can have as many as 16 copies. A single database availability group can have up to 16 mailbox servers that provide automatic database-level recovery. Any server in

a database availability group can host a copy of a mailbox database from any other server in the database availability group.

This seamless high-availability functionality is possible because mailbox databases are disconnected from servers and the same globally unique identifier (GUID) is assigned to every copy of a mailbox database. Because there are no storage groups, continuous replication occurs at the database level. Transaction logs are replicated to each member of a database availability group that has a copy of a mailbox database and are replayed into the copy of the mailbox database. Failover can occur at either the database level or the server level.

With regard to architecture, Exchange Server 2016 is significantly different from early releases of Exchange and furthers the server role consolidation that began with Exchange 2013. While Exchange 2010 and other earlier releases had components split into different server roles for scaling out Exchange organizations, Exchange 2016 does not have separate server roles for Hub Transport servers, Unified Messaging servers or even Client Access servers. As the related components are all now part of the Mailbox Server role, there are no longer any other roles for internal Exchange servers. That said, although Exchange Server 2013 did not include the Edge Transport role until the release of Service Pack 1, the Edge Transport role also is included with Exchange 2016 and available for deployment in perimeter zones. This means you can use Edge Transport servers running this latest version of Exchange to add a layer of security to the Exchange organization.

These architecture changes mean that Mailbox servers storing the active database copy for a mailbox perform all the data processing, data rendering, and data transformation required. The Mailbox server connects the client and performs authentication. Supported protocols for client connections include HTTP, POP, IMAP, RPC over HTTP, MAPI over HTTP and SMTP. As RPC is no longer supported as a direct access protocol, all Outlook client connections must take place using either RPC over HTTP or MAPI over HTTP, with the latter being preferred.

It's important to point out that Exchange 2016 is designed to work with Outlook 2010 and higher and also continues to support Outlook Web App for mobile access. Rather than connecting to servers using Fully Qualified Domain Names as was done in the past, Outlook 2010 and higher use Autodiscover to create connection points based on the domain portion of the user's primary SMTP address and each mailbox's Globally Unique Identifier (GUID).

The simplified architecture reduces the namespace requirements for Exchange site designs. If you're coexisting with Exchange 2010 or you're installing a new Exchange 2016 organization, you need only one namespace for client protocols and one namespace for Autodiscover. To continue to support SMTP, you also need an SMTP namespace.

REAL WORLD As discussed in detail in Chapter 25 "Optimizing Web and Mobile Access," Outlook Web App (OWA) is a browser-based application that is

accessed via IIS running on your Mailbox servers. Outlook Web App is also referred to as Outlook on the Web in some documentation from Microsoft and in some parts of the Exchange UI. Outlook on the Web is the new name and branding identity for Outlook Web App. For consistency, this text references Outlook Web App or OWA, which is what you'll see in most parts of the UI until re-branding is complete.

Selecting Hardware for Exchange 2016

Before you deploy Exchange Server 2016, you should carefully plan the messaging architecture. As part of your implementation planning, you need to look closely at preinstallation requirements and the hardware you will use. Exchange Server is a complex messaging platform with many components that work together to provide a comprehensive solution for routing, delivering, and accessing email messages, voice-mail messages, faxes, contacts, and calendar information.

If you're using Exchange Online, Microsoft provides the server hardware. Otherwise, for on-premises implementations, Exchange Server 2016 should run on a system with adequate memory, processing speed, and disk space. You also need an appropriate data-protection and system-protection plan at the hardware level.

Exchange Server 2016 requires two different types of server hardware. You want to select hardware for Mailbox servers with scaling *up* in mind while selecting hardware for Edge Transport servers with scaling *out* in mind. Scaling up typically means adding additional or faster, better CPUs and memory to existing servers to meet capacity needs. Scaling out typically means adding additional servers to meet capacity needs.

Key guidelines for choosing hardware for Exchange Server are as follows:

- **Memory** The minimum random access memory (RAM) is 8 gigabytes (GB) for servers with the Mailbox Server role and 4 GB for Edge Transport servers. In most cases, you'll want to have at least twice the recommended minimum amount of memory. The primary reason for this is performance. Most of the Mailbox server installations I run use 16 GB of RAM as a starting point, even in small installations. In multiple Exchange server installations, the Mailbox server should have at least 2 GB of RAM plus 5 megabytes (MB) of RAM per mailbox (with a minimum of 8 GB regardless). For all Exchange server configurations, the paging file should be at least equal to the amount of RAM in the server plus 10 MB, with a maximum size of 32778 MB if the server has more than 32 GB of RAM.
- **CPU** Exchange Server 2016 runs on the x64 family of processors from AMD and Intel, including AMD64 and Intel 64. You can achieve significant performance improvements with a high level of processor cache. Look closely at the L1, L2, and L3 cache options available—a higher cache can yield much better performance overall. Look also at the speed of the front-side bus. The faster the bus speed, the faster the CPU can access memory.
- **SMP** Exchange Server 2016 supports symmetric multiprocessors, and you'll see significant performance improvements if you use multiple CPUs—not just multiple cores in a single CPU. Although the clock speed of the CPU is important, so are the number of logical processor cores and the number of threads that can be simultaneously processed. That said, if Exchange Server is supporting a small organization with a single domain, one CPU with multiple cores may be enough. If the server supports a medium or large organization or handles mail for multiple domains,

you will want to consider adding processors. When it comes to processor cores, I prefer two multicore processors to a single processor with the same number of cores, given current price and performance tradeoffs. As part of the preferred architecture, Microsoft recommends dual socket servers, with 16 to 24 cores each for large organizations.

- **Disk drives** The data storage capacity you need depends entirely on the number and size of the data that will pass through, be journaled on, or stored on the Exchange server. You need enough disk space to store all data and logs, plus workspace, system files, and virtual memory. Input/output (I/O) throughput is just as important as drive capacity. Rather than use one large drive, you should use several drives, which allows you to configure fault tolerance. As part of your hardware planning, it's important to point out that Exchange 2016 supports multiple databases on the same volume, allowing you to have a mix of active and passive copies on a single volume. Keep in mind, however, the input/output per second (IOPS) capabilities for the underlying physical disks. Also note that even if you've been assigned multiple logical unit numbers (LUNs) for use from storage these different LUNs may be spread over the same physical disks. Make sure your hardware uses battery-backed write cache controllers as this will help prevent data loss during an unexpected loss of power. Finally, serially attached SCSI (SAS) disks are preferred to SATA disks as they provide better IO and a lower failure rate.
- **File formats** Although Windows Server supports several file systems, only the NTFS and ReFS file system should be used with Exchange Server 2016. Use NTFS for the system partition and for any partition that stores Exchange binary files or files generated by diagnostic logging. Use NTFS or ReFS for partitions containing database files, content indexing files and transaction log files. As discussed in more detail in Chapter 14, "Preparing for Exchange 2016," ReFS is preferred to NTFS as it is a more efficient and more resilient filesystem. If you use ReFS with Exchange, be sure to disable the integrity feature.
- **Data protection** You can add protection against unexpected drive failures by using redundant storage. For the boot and system disks, use mirroring on internal drives. However, when you use availability groups and other high-availability features of Exchange Server, you might not want to use redundant storage for Exchange data and logs. You also might not want to use expensive disk storage systems either. Instead, deploy multiple Exchange servers in each of your database availability groups. As part of the preferred architecture, Microsoft recommends using disk mirroring (RAID 1) for the operating system volume with the rest of the storage configured as Just a Bunch of Disks (JBOD). This preferred architecture assumes that you have deployed multiple mailbox servers in a highly available configuration.
- **Data encryption** Consider using BitLocker to encrypt the operating system volume and Exchange data volumes. Data encryption reduces concerns about data theft for those with physical access to the server. If you deploy highly available Exchange servers without BitLocker and later want to add BitLocker, be sure to put the server you are working with in maintenance mode first.
- **Uninterruptible power supply** Exchange Server 2016 is designed to maintain

database integrity at all times and can recover information using transaction logs. This doesn't protect the server hardware, however, from sudden power loss or power spikes, both of which can seriously damage hardware. To prevent this, connect your server to an uninterruptible power supply (UPS). A UPS gives you time to shut down the server or servers properly in the event of a power outage. Proper shutdown is especially important on servers using write-back caching controllers. These controllers temporarily store data in cache. Without proper shutdown, this data can be lost before it is written to disk. To prevent data loss, write-back caching controllers typically have batteries that help ensure that changes can be written to disk after the system comes back online.

If you follow these hardware guidelines, you'll be well on your way to success with Exchange Server 2016. It's important to note that beginning with Windows Server 2012, dynamic disks are being phased out in favor of Storage Spaces. However, for mirroring boot and system volumes on internal disks, Microsoft recommends continuing to use dynamic disks and RAID 1.

If you decide to use software-based redundant storage, remember that storage arrays typically already have an underlying redundant storage configuration and you might have to use a storage array-specific tool to help you distinguish between LUNs and the underlying physical disks. Herein, I focus on software-based redundancy implemented with RAID or Storage Spaces rather than the underlying hardware redundancy implemented in storage arrays.

Windows Server is transitioning to standards-based storage beginning with Windows Server 2012. This transition means several popular tools and favored features are being phased out. Officially, a tool or feature that is being phased out is referred to as *deprecated*. When Microsoft deprecates a tool or feature, it might not be in future releases of the operating system (while continuing to be available in current releases). Rather than not cover popular tools and features, I've chosen to discuss what is actually available in the current operating system, including both favored standbys and newer options. One of these newer options is Storage Spaces. With Storage Spaces:

- Simple volumes can stretch across multiple disks, similar to disk striping with parity (RAID 0).
- Mirrored volumes are mirrored across multiple disks. Although this is similar to disk mirroring (RAID 1), it is more sophisticated in that data is mirrored onto two or three disks at a time. If a storage space has two or three disks, you are fully protected against a single disk failure, and if a storage space has five or more disks, you are fully protected against two simultaneous disk failures.
- Parity volumes use disk striping with parity. Although this is similar to RAID 5, it is more sophisticated in that there are more protections and efficiencies.

Navigating Exchange 2016 Editions

Several editions of Exchange Server 2016 are available, including Exchange Server 2016 Standard and Exchange Server 2016 Enterprise. The various server editions support the same core features and administration tools, which means you can use the techniques discussed throughout this book regardless of which Exchange Server 2016 edition you are using. For reference, the specific feature differences between Standard Edition and Enterprise Edition are as follows:

- **Exchange Server 2016 Standard** Designed to provide essential messaging services for small to medium organizations and branch office locations. This server edition supports up to five databases.
- **Exchange Server 2016 Enterprise** Designed to provide essential messaging services for organizations with increased availability, reliability, and manageability needs. This server edition supports up to 100 databases (including all active databases and copies of databases) on a particular server. It's important to note that this is a substantial reduction in the number of supported databases as compared to Exchange 2010.

NOTE Throughout this book, I refer to Exchange Server 2016 in different ways, and each has a different meaning. Typically, I refer to the software product as *Exchange 2016* or as *Exchange Server*, which you can take to mean *Microsoft Exchange Server 2016*. When necessary, I use *Exchange Server 2016* to draw attention to the fact that I am discussing a feature that's new or has changed in the most recent version of the product. Each of these terms means essentially the same thing. If I refer to a previous version of Exchange Server, I always do so specifically, such as Exchange 2010 or Exchange 2013. Finally, I often use the term *Exchange server* (note the lowercase *s* in server) to refer to an actual server computer, as in "There are eight Exchange servers in this database availability group."

REAL WORLD Microsoft provides a single binary for x64 systems, and the same binary file is used for both the Standard and Enterprise editions. The license key provided during installation is what determines which edition is established during installation.

You can use a valid product key to upgrade from a trial edition to the Standard edition or the Enterprise edition of Exchange Server 2016 without having to reinstall. Using a valid product key, you can also upgrade from the Standard to the Enterprise edition. You can also relicense an Exchange server by entering a new product key for the installed edition, which is useful if you accidentally used the same product key on multiple servers and want to correct the mistake.

There are several caveats. When you change the product key on a Mailbox server, you must restart the Microsoft Exchange Information Store service to apply the

change. Additionally, you cannot use product keys to downgrade editions. To downgrade editions, you must uninstall Exchange Server and then reinstall Exchange Server.

Exchange 2016 is supported on Windows Server 2012, Windows Server 2012 R2 and Windows Server 2016. You can install Exchange Server 2016 on servers running full-server installations of Windows Server 2012 Standard, Windows Server 2012 Datacenter or later. You cannot install Exchange 2016 on servers running server core or minimal server interface. With Windows Server 2012 or Windows Server 2012 R2, you must convert the server core or minimal server interface installation to a full installation by running the following command from an elevated PowerShell prompt:

```
Install-WindowsFeature Server-Gui-Mgmt-Infra, Server-Gui-Shell -Restart
```

In addition to being able to deploy Exchange 2016 on physical servers, you can deploy Exchange 2016 in a virtualized environment on any version of Windows Server with Hyper-V technology or Microsoft Hyper-V Server as well as other hypervisors that have been validated under the Windows Server Virtualization Validation Program. If you do so, keep in mind that making virtual machine snapshots of an Exchange guest virtual machine isn't supported, as the state of Exchange 2016 would not be fully preserved.

REAL WORLD There are many arguments for and against using virtualization with Exchange Server. Typically, Microsoft recommends using physical servers as part of the preferred architecture, especially for large organizations. The primary reason for this is that virtualization adds complexity while increasing the management overhead as well.

A client accessing an Exchange server requires a Client Access License (CAL). With either Exchange Server edition, the client can use a Standard CAL, an Enterprise CAL, or both. The Standard CAL allows for the use of email, shared calendaring, contacts, task management, Microsoft Outlook Web App (OWA), and Exchange ActiveSync. The Enterprise CAL allows for the use of unified messaging, advanced mobile management, data loss prevention, and custom retention policies. An Enterprise CAL is sold as an add-on to the Standard CAL. A client must have one Standard CAL and one Enterprise CAL add-on to make full use of all Exchange Server features.

Beyond the editions and CALs, Exchange Server 2016 has several variants. Microsoft offers on-premises and online implementations of Exchange Server. An on-premises Exchange Server is one that you install in your organization. An online Exchange Server is delivered as a subscription service from Microsoft. In Exchange Server 2016, you can manage both on-premises and online implementations of Exchange Server using the same management tools. These implementations can be separate from each other, or you can configure a hybrid installation that allows single sign-on and easy movement of mailboxes and databases between on-premises and online implementations.

As a prerequisite for installing any server running any on-premises version of Exchange Server 2016, Active Directory must be at Windows Server 2008 forest functionality mode or higher. Additionally, all domain controllers in the Active Directory forest must

be running Windows Server 2008 Standard, Enterprise or Datacenter, Windows Server 2012 Standard or Datacenter, or later editions of Windows Server.

NOTE Using Active Directory with Exchange Server 2016 is covered in more detail in the “Using Exchange 2016 with Active Directory” section of this chapter and the “Integrating Exchange with Active Directory” section of Chapter 14.

Additionally, Exchange Server 2016 supports IPv6 only when IPv4 is also installed and enabled. When you deploy IPv6, Exchange servers can send data to and receive data from devices, clients, and servers that use IPv6 addresses. You install Exchange 2016 using Exchange Setup. Exchange 2016 requires:

- Microsoft .NET Framework version 4.5.2 (<http://go.microsoft.com/fwlink/p/?linkid=518380>)
- Microsoft Unified Communications Managed API 4.0, Core Runtime 64-bit (<http://go.microsoft.com/fwlink/p/?linkid=258269>)
- Windows Management Framework 4.0 (<https://www.microsoft.com/en-us/download/details.aspx?id=40855>).

If you don't install these additional components prior to running Exchange Setup, the Readiness Checks will fail and links to these resources will be provided. If this happens, you can use the links provided to obtain and install the components and then simply click Retry to have Setup perform the readiness checks again. Once these checks pass, you'll be able to continue with the installation.

With Edge Transport servers, you must install the Active Directory Lightweight Directory Service (AD LDS) feature of Windows. To install, use the following command:

```
Install-WindowsFeature ADLDS
```

With Mailbox servers, you must install a number of Windows features prior to installing Exchange Server. To install these features, use the following command:

```
Install-WindowsFeature AS-HTTP-Activation, Desktop-Experience, NET-Framework-45-Features, RPC-over-HTTP-proxy, RSAT-Clustering, RSAT-Clustering-CmdInterface, RSAT-Clustering-Mgmt, RSAT-Clustering-PowerShell, Web-Mgmt-Console, WAS-Process-Model, Web-Asp-Net45, Web-Basic-Auth, Web-Client-Auth, Web-Digest-Auth, Web-Dir-Browsing, Web-Dyn-Compression, Web-Http-Errors, Web-Http-Logging, Web-Http-Redirect, Web-Http-Tracing, Web-ISAPI-Ext, Web-ISAPI-Filter, Web-Lgcy-Mgmt-Console, Web-Metabase, Web-Mgmt-Console, Web-Mgmt-Service, Web-Net-Ext45, Web-Request-Monitor, Web-Server, Web-Stat-Compression, Web-Static-Content, Web-Windows-Auth, Web-WMI, Windows-Identity-Foundation
```

If you don't install these Windows features prior to running Exchange Setup, the setup process will try to install these components for you. Some of these components require a server restart. Others may not install completely until you exit Exchange Setup and repeat the Exchange setup process. Either way, you'll likely need to repeat the Exchange setup process. Rather than trying to type the long command from this text, see the

companion website for this book at <http://www.williamrstanek.com/exchangeserver/> for command text you can copy and paste.

Exchange 2016 has a different set of management tools than its predecessors. When you install a Mailbox server or an Edge Transport server, the management tools are installed automatically. You can use Exchange Setup to install the management tools on domain-joined computers running 64-bit editions of Windows 8.1, Windows 10 as well as Windows Server 2012 or later. Although Exchange 2016 has management tools that can be installed, you can perform most management actions remotely using a standard web browser and Windows PowerShell, and you'll learn more about this later in this chapter.

Exchange Server 2016 uses the Windows Installer (the Installer) and has a fully integrated installation process. This means you can configure Exchange Server 2016 much like you can any other application you install on the operating system. The installation can be performed from a command prompt as well.

Chapter 14 provides detailed instructions for installing Exchange Server 2016. You install Exchange 2016 only on domain-joined computers. Whether you use the Standard or Enterprise edition, you have similar options. You can install an internal messaging server by selecting the Mailbox role or a perimeter zone server by selecting the Edge Transport role. Generally, you will not want an internal Exchange server to also be configured as a domain controller with a global catalog.

When you start an installation, Setup checks the system configuration to determine the local time zone, the operating system, the logged-on user, and the status of the registry keys related to Exchange Server 2016. Installation will fail if you are trying to run Setup on an operating system that isn't supported.

After checking the system configuration, Setup allows you to check for updates to the installation process, provided the server has a connection to the Internet. Setup then checks available space on the %SystemDrive% to ensure a temporary folder under %SystemDrive%\Windows\Temp\ExchangeSetup can be used during the installation process. Before you install Exchange 2016, you should ensure the drive on which you plan to install Exchange has at least 32 GB of disk space.

When done copying its work files to the temporary folder, Setup tries to connect to a domain controller and validate the state of Active Directory. If Setup cannot find a domain controller or encounters other errors when validating Active Directory, the installation process will fail and you'll see related errors during the readiness checks.

IMPORTANT By default, Setup chooses a domain controller in the local domain and site. In order to determine the domain information and contact a domain controller, the computer on which you are installing Exchange 2016 must be domain joined and have properly configured TCP/IP settings, and DNS name resolution must be properly configured in your organization. Because Active Directory site configuration also is important for installing Exchange 2016 and setting up an Exchange organization, ensure Active Directory sites and subnets are

properly configured prior to installing Exchange 2016.

Once connected to a domain controller, Setup selects a global catalog server to work with and then looks for an Exchange Configuration container within Active Directory. Setup next determines the organization-level operations that need to be performed, which can include initializing Active Directory, updating Active Directory schema, establishing or updating the Exchange organization configuration, and updating the domain configuration.

As you continue through Setup, you'll be able to select the server roles to install, the install location, and more. With the exception of the working files, which are copied to the temporary folder, no changes are made until the server passes the readiness checks. Normally, even when problems are encountered, Setup will continue all the way to the readiness checks. As part of the readiness checks, Setup checks for required components, such as those listed previously.

Other required components include Windows Features that Setup will install automatically if they aren't already installed. These features include Desktop Experience, many components of IIS, Windows Identity Foundation, and the administrative tools for clustering. Although you can manually install these features, it's a long list; Setup will do the work for you if you let it and don't mind having to repeat the setup process.

Exchange 2016 includes the following anti-spam capabilities:

- **Sender filtering** Allows administrators to maintain a list of senders who are blocked from sending messages to the organization. Administrators can block individual senders by email address. Administrators also can block all senders from domains and subdomains.
- **Recipient filtering** Allows administrators to block message delivery to nonexistent recipients, distribution lists for internal users only, and mailboxes for internal use only. Exchange performs recipient lookups on incoming messages and block messages, which prevents certain types of attacks and malicious attempts at information discovery.
- **Sender ID verification** Verifies that incoming email messages are from the Internet domain from which they claim to come. Exchange verifies the sender ID by examining the sender's IP address and comparing it to the related security record on the sender's public DNS server.
- **Content filtering** Uses intelligent message filtering to scan message content and identify spam. Spam can be automatically deleted, quarantined, or filed as junk email.

TIP Using the Exchange Server management tools, administrators can manage messages sent to the quarantine mailbox and take appropriate actions, such as deleting messages, flagging them as false positives, or allowing them to be delivered as junk email. Messages delivered as junk email are converted to plain text to strip out any potential viruses they might contain.

- **Sender reputation scoring** Helps to determine the relative trustworthiness of unknown senders through sender ID verification and by examining message content and sender behavior history. A sender can then be added temporarily to the Blocked Senders list.

The way you use these features will depend on the configuration of your Exchange organization. If you've deployed Edge Transport servers, you enable and configure these features on your Edge Transport servers. Otherwise, you enable and configure these features on your Mailbox servers.

Exchange 2016 also has anti-malware capabilities, which are enabled by default. Malware scanning is performed on all messages at the server level as messages are sent or received. When users open and read messages in their mailboxes, the messages they see have already been scanned. Exchange Server checks for updates to malware definitions every hour. Exchange downloads the malware engines and definitions using a TCP connection over port 80 from the Internet.

TIP Normally, you'll manually perform the first download of the anti-malware engine and definition updates prior to placing a server into production so you can verify that the initial process was successful and then configure default anti-malware policy prior to users having access to a server.

Although these anti-spam and anti-malware features are extensive, they are not comprehensive. For comprehensive protection, you can pair these features with a cloud-based service, such as Microsoft Exchange Online Protection. By combining the built-in anti-spam and anti-malware features with a cloud-based protection service you can set up substantial, layered protection. Additionally, if you use a third-party anti-malware solution for Exchange 2016, you can disable the built-in anti-malware filtering.

Using Exchange 2016 with Windows Server

When you install Exchange Server on a server operating system, Exchange Server makes extensive modifications to the environment. These modifications include additional system services, integrated authentication, and additional security groups.

Services for Exchange Server

When you install Exchange Server and Forefront Protection for Exchange Server on Windows, multiple services are installed and configured on the server. A summary of key services, how they are used, and which server components they are associated with follows:

- **IIS Admin** Enables the server to administer the IIS metabase. The IIS metabase stores configuration information for web applications used by Exchange. Exchange servers need IIS for WinRM and remote Powershell. Mailbox servers need IIS for OWA and Web services.
- **Microsoft Exchange Active Directory Topology** Provides Active Directory topology information to Exchange services. If this service is stopped, most Exchange services will not be able to start.
- **Microsoft Exchange Anti-Spam Update** Maintains the anti-spam data for Forefront Protection on an Exchange server.
- **Microsoft Exchange EdgeSync** Provides EdgeSync services between Mailbox and Edge servers.
- **Microsoft Exchange Frontend Transport** Proxies inbound and outbound SMTP connections.
- **Microsoft Exchange IMAP4 Backend** Provides IMAP4 services to mailboxes.
- **Microsoft Exchange IMAP4** Provides IMAP4 services to clients.
- **Microsoft Exchange Information Store** Manages the Microsoft Exchange Information Store. This includes mailbox stores and public folder stores.
- **Microsoft Exchange Mailbox Assistants** Manages assistants responsible for calendar updates, booking resources, and other mailbox processing.
- **Microsoft Exchange Mailbox Replication** Enables online mailbox moves by processing mailbox move requests.
- **Microsoft Exchange Mailbox Transport Delivery** Receives mail items from the Transport service and ensures they are processed and then delivered into mailbox.
- **Microsoft Exchange Mailbox Transport Submission** Receives mail items being sent and ensures they are converted from MAPI to MIME and then submitted to the Transport service.
- **Microsoft Exchange POP3 Backend** Provides Post Office Protocol version 3 (POP3) services to mailboxes.
- **Microsoft Exchange POP3** Provides Post Office Protocol version 3 (POP3) services to clients.
- **Microsoft Exchange Replication** Provides replication functionality used for

continuous replication.

- **Microsoft Exchange RPC Client Access** Manages client remote procedure call (RPC) connections for Exchange Server.
- **Microsoft Exchange Search** Handles queries and controls indexing of mailboxes to improve search performance.
- **Microsoft Exchange Server Extension for Windows Server Backup** Provides extensions for Windows Server Backup that allow you to back up and recover Exchange application data using Windows Server Backup.
- **Microsoft Exchange Service Host** Provides a host for essential Exchange services.
- **Microsoft Exchange Throttling** Provides throttling functions to limit the rate of user operations.
- **Microsoft Exchange Transport Log Search** Provides search capability for Exchange transport log files.
- **Microsoft Exchange Transport** Provides mail transport for Exchange Server.
- **Microsoft Exchange Unified Messaging Call Router** Provides capabilities necessary for routing calls.
- **Microsoft Exchange Unified Messaging** Enables voice and fax messages to be stored in Exchange and gives users telephone access to email, voice mail, the calendar, contacts, or an automated attendant.
- **Secure Socket Tunneling Protocol Service** Provides support for Secure Socket Tunneling Protocol (SSTP) for securely connecting to remote computers.
- **Web Management Service** Enables remote and delegated management for the web server, sites, and applications.
- **Windows Remote Management** Implements the WS-Management protocol. Required for remote management using the Exchange console and Windows PowerShell.
- **World Wide Web Publishing Service** Provides web connectivity and administration features for IIS.

Exchange Server Authentication and Security

In Exchange Server 2016, email addresses, distribution groups, and other directory resources are stored in the directory database provided by Active Directory. Active Directory is a directory service running on Windows domain controllers. When there are multiple domain controllers, the controllers automatically replicate directory data with each other using a multimaster replication model. This model allows any domain controller to process directory changes and then replicate those changes to other domain controllers.

The first time you install Exchange Server 2016 in a Windows domain, the installation process updates and extends Active Directory to include objects and attributes used by Exchange Server 2016. Unlike earlier releases of Exchange Server, you do not use Active Directory Users And Computers to manage mailboxes, messaging features, messaging options, or email addresses associated with user accounts. You perform these

tasks using the Exchange Management tools.

Exchange Server 2016 fully supports the Windows Server security model and by default relies on this security mechanism to control access to directory resources. This means you can control access to mailboxes and membership in distribution groups and you can perform other Exchange security administration tasks through the standard Windows Server permission set. For example, to add a user to a distribution group, you simply make the user a member of the distribution group in Active Directory Users And Computers.

Because Exchange Server uses Windows Server security, you can't create a mailbox without first creating a user account that will use the mailbox. Every Exchange mailbox must be associated with a domain account—even those used by Exchange for general messaging tasks. In the Exchange Admin Center, you can create a new user account as part of the process of creating a new mailbox.

You manage Exchange servers according to their roles and the type of information you want to manage using the Exchange Admin Center. You'll learn more about this in Chapter 16, "Exchange 2016 Administration Essentials."

Exchange Server Security Groups

Exchange Server 2016 uses predefined universal security groups to separate administration of Exchange permissions from administration of other permissions. When you add an administrator to one of these security groups, the administrator inherits the permissions permitted by that role.

The predefined security groups have permissions to manage the following types of Exchange data in Active Directory:

- **Organization configuration data** This type of data is not associated with a specific server and is used to manage databases, policies, address lists, and other types of organizational configuration details.
- **Server configuration data** This type of data is associated with a specific server and is used to manage the server's messaging configuration.
- **Recipient configuration data** This type of data is associated with mailboxes, mail-enabled contacts, and distribution groups.

The predefined groups are as follows:

- **Compliance Management** Members of this group have permission to configure compliance settings.
- **Delegated Setup** Members of this group have permission to install and uninstall Exchange on provisioned servers.
- **Discovery Management** Members of this group can perform mailbox searches for data that meets specific criteria.
- **Exchange Servers** Members of this group are Exchange servers in the organization. This group allows Exchange servers to work together.

- **Exchange Trusted Subsystem** Members of this group are Exchange servers that run Exchange cmdlets using WinRM. Members of this group have permission to read and modify all Exchange configuration settings as well as user accounts and groups.
- **Exchange Windows Permissions** Members of this group are Exchange servers that run Exchange cmdlets using WinRM. Members of this group have permission to read and modify user accounts and groups.
- **Help Desk** Members of this group can view any property or object within the Exchange organization and have limited management permissions, including the right to change and reset passwords.
- **Hygiene Management** Members of this group can manage the anti-spam and antivirus features of Exchange.
- **Managed Availability Servers** Every Exchange 2016 server is a member of this group. Managed availability is an internal process that provides native health monitoring and recovery for protocol processes to ensure availability of Exchange services. For more information, see Chapter 18, “Implementing Availability Groups.”
- **Organization Management** Members of this group have full access to all Exchange properties and objects in the Exchange organization.
- **Public Folder Management** Members of this group can manage public folders and perform most public folder management operations.
- **Recipient Management** Members of this group have permissions to modify Exchange user attributes in Active Directory and perform most mailbox operations.
- **Records Management** Members of this group can manage compliance features, including retention policies, message classifications, and transport rules.
- **Server Management** Members of this group can manage all Exchange servers in the organization but do not have permission to perform global operations.
- **UM Management** Members of this group can manage all aspects of unified messaging, including unified messaging server configuration and unified messaging recipient configuration.
- **View-Only Organization Management** Members of this group have read-only access to the entire Exchange organization tree in the Active Directory configuration container and read-only access to all the Windows domain containers that have Exchange recipients.

Using Exchange 2016 with Active Directory

Exchange Server 2016 is tightly integrated with Active Directory. Not only does Exchange Server 2016 store information in Active Directory, but it also uses the Active Directory routing topology to determine how to route messages within the organization. Routing to and from the organization is handled using transport servers.

Understanding How Exchange Stores Information

Exchange stores four types of data in Active Directory: schema data (stored in the Schema partition), configuration data (stored in the Configuration partition), domain data (stored in the Domain partition), and application data (stored in application-specific partitions). In Active Directory, schema rules determine what types of objects are available and what attributes those objects have. When you install the first Exchange server in the forest, the Active Directory preparation process adds many Exchange-specific object classes and attributes to the schema partition in Active Directory. This allows Exchange-specific objects, such as agents and connectors, to be created. It also allows you to extend existing objects, such as users and groups, with new attributes, such as attributes that allow user objects to be used for sending and receiving email. Every domain controller and global catalog server in the organization has a complete copy of the Schema partition.

During the installation of the first Exchange server in the forest, Exchange configuration information is generated and stored in Active Directory. Exchange configuration information, like other configuration information, is also stored in the Configuration partition. For Active Directory, the configuration information describes the structure of the directory, and the Configuration container includes all of the domains, trees, and forests, as well as the locations of domain controllers and global catalogs. For Exchange, the configuration information is used to describe the structure of the Exchange organization. The Configuration container includes lists of templates, policies, and other global organization-level details. Every domain controller and global catalog server in the organization has a complete copy of the Configuration partition.

In Active Directory, the Domain partition stores domain-specific objects, such as users and groups, and the stored values of attributes associated with those objects. As you create, modify, or delete objects, Exchange stores the details about those objects in the Domain partition. During the installation of the first Exchange server in the forest, Exchange objects are created in the current domain. Whenever you create new recipients or modify Exchange details, the related changes are reflected in the Domain partition as well. Every domain controller has a complete copy of the Domain partition for the domain for which it is authoritative. Every global catalog server in the forest maintains information about a subset of every Domain partition in the forest.

Understanding How Exchange Routes Messages

Within the organization, the Transport service on Mailbox servers use the information about sites stored in Active Directory to determine how to route messages, and they can also route messages across site links. These servers do this by querying Active Directory about its site membership and the site membership of other servers, and then using the information they discover to route messages appropriately. Because of this, when you are deploying an Exchange Server 2016 organization, no additional configuration is required to establish routing in the Active Directory forest.

For mail delivery within the organization, additional routing configuration is necessary only in these specific scenarios:

- If you deploy an Exchange Server 2016 organization with multiple forests, you must install Exchange Server 2016 in each forest and then connect the forests using appropriate cross-forest trusts. The trust allows users to see address and availability data across the forests.
- In an Exchange Server 2016 organization, if you want direct mail flow between Exchange servers in different forests, you must configure SMTP send connectors and SMTP receive connectors on the Mailbox servers that should communicate directly with each other.

You can use two types of Mail Transport servers: Mailbox servers and Edge Transport servers. You deploy Mailbox servers within the organization. The Transport service on Mailbox servers handles mail delivery and receipt of mail. Two services are used to deliver mail items to, and receive mail items from, other servers:

- **Microsoft Exchange Mailbox Transport Delivery service** Handles inbound mail items. After receiving mail items for delivery to a mailbox on the current server, the service submits the mail items for processing and then delivers them into the appropriate mailbox database on the server.
- **Microsoft Exchange Mailbox Transport Submission service** Handles outbound mail items. After receiving mail items for submission, the service ensures messages are converted from MAPI to MIME and then passes them along to the Transport service. The Transport service then routes the mail items for delivery.

With Mailbox servers as your transports, no other special configuration is needed for message routing to external destinations. You must configure only the standard mail setup, which includes identifying DNS servers to use for lookups. With Edge Transport servers, you can optimize mail routing and delivery by configuring one-way synchronization from the internal Mailbox servers to the perimeter network's Edge Transport servers. Beyond this, no other special configuration is required for mail routing and delivery.

You deploy Edge Transport servers in the organization's perimeter network for added security. Typically, a perimeter network is a secure network set up outside the organization's private network. When you have Edge Transport servers, mail items from outside the organization are received first by the Edge transport servers, which can perform anti-malware and anti-spam checks before passing along mail items to internal

Mailbox servers for delivery. Mail items for submission outside the organization are passed from internal Mailbox servers to Edge Transport servers which then submit the mail items for delivery outside the organization.

Additional Tools and Options

The Exchange Toolbox is the only remaining MMC-based console for Exchange Server administration. The Toolbox, shown in Figure 13-1, provides access to a suite of related management tools, including:

- **Details Templates Editor** Helps administrators customize client-side GUI presentation of object properties accessed through address lists. You can use this tool to customize the presentation of contacts, users, groups, public folders, and more in the client interface.
- **Remote Connectivity Analyzer** Allows administrators to perform connectivity tests for inbound email, ActiveSync, Exchange Web Services, Outlook Anywhere, and MAPI over HTTP.
- **Queue Viewer** Allows administrators to track message queues and mail flow. Also allows administrators to manage message queuing and remove messages.

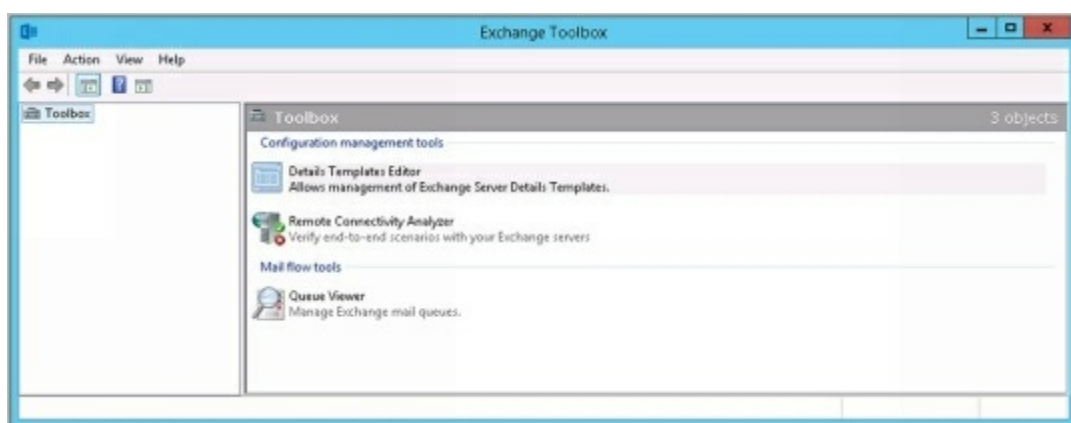


FIGURE 13-1 The Exchange Toolbox.

On any computer where you've installed the Exchange management tools, you can access the Exchange Toolbox from Start. Whether you are working with the Start menu or the Start screen, you can pin the Exchange Toolbox to the desktop taskbar by right-clicking the related icon and then selecting Pin To Taskbar.

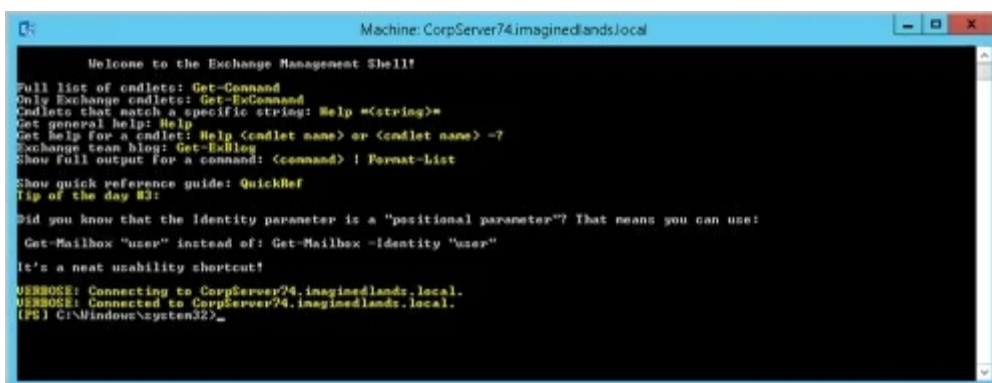
Other tools that you might want to use with Exchange Server include:

- **DNS console** Manages the DNS service.
- **Event Viewer** Manages events and logs.
- **Failover Cluster Management** The Failover Cluster Management tools and the related command-line interface must be installed on your Exchange 2016 servers. This allows you to use scripts for managing availability groups.
- **IIS Manager** Manages Web servers used by Exchange as well as the management service configuration.
- **Server Manager** Provides setup and configuration options for the local server as well as options for managing roles, features and related settings on remote servers.

You access most of these tools from the Tools menu in Server Manager. Server Manager can be started by clicking the Server Manager icon in the taskbar. With Windows Server 2012 and later, you also can start Server Manager by typing **Server Manager** in the Search box and pressing **Enter**.

Although the graphical tools provide just about everything you need to work with Exchange organizations, there are many times when you might want to work from the command line, especially if you want to automate installation, administration, or maintenance with scripts. To help with all your command-line needs, Exchange Server includes Exchange Management Shell.

Exchange Management Shell is an extension shell for Windows PowerShell that includes a wide array of built-in commands for working with Exchange Server. On any computer where you've installed the Exchange management tools, you'll be able to access Exchange Management Shell from Start. Whether you are working with the Start menu or the Start screen, you can pin Exchange Management Shell to the desktop taskbar by right-clicking the related icon and then selecting Pin To Taskbar. Exchange Management Shell is shown in Figure 13-2.



```
Machine: CorpServer74.imagedlands.local

Welcome to the Exchange Management Shell!

Full list of cmdlets: Get-Command
Only Exchange cmdlets: Get-ExCommand
Cmdlets that match a specific string: Help <string>
Get general help: Help
Get help for a cmdlet: Help <cmdlet name> or <cmdlet name> -?
Exchange team blog: Get-ExBlog
Show full output for a cmdlet: <command> | Format-List
Show quick reference guide: QuickRef
Tip of the day #3:
Did you know that the Identity parameter is a "positional parameter"? That means you can use:
Get-Mailbox "user" instead of: Get-Mailbox -Identity "user"
It's a neat usability shortcut!
VERBOSE: Connecting to CorpServer74.imagedlands.local.
VERBOSE: Connected to CorpServer74.imagedlands.local.
(PS) C:\Windows\system32>
```

FIGURE 13-2 Exchange Management Shell.

REAL WORLD Exchange Admin Center is a web-based management console that runs as an application on your Mailbox servers. When you install the Mailbox server role for Exchange 2016, the server is configured automatically with a Windows PowerShell gateway that acts as a proxy service. This proxy service allows you to run remote commands in web browsers and in remote sessions. Whenever you work with Exchange Admin Center or Exchange Management Shell, the commands are executed via this proxy—even if you log on locally. Thus, every time you work with Exchange Server, you are using a remote session.

When you log in to the Exchange Admin Center, you are using the Default Web Site running on Internet Information Services (IIS) which processes your actions. Every command you perform in Exchange Admin Center is remotely executed via the Windows PowerShell gateway, as is any command you perform in Exchange Management Shell. Any task you can perform in Exchange Admin Center can be performed in Exchange Management Shell.

Exchange Management Shell is designed to be run only on domain-joined computers.

The basics of working with Exchange Management Shell are straightforward:

- Type **get-command** to get a full list of all available cmdlets on the server.
- Type **get-excommand** to get a full list of all Exchange-specific cmdlets available.
- Type **help cmdletName** to get help information, where *cmdletName* is the name of the command you are looking up.

IMPORTANT When you are working with Exchange Management Shell, the default recipient scope is set the same as your logon domain. If you are in multi-domain environment and want to work with recipients throughout the Active Directory forest, make sure the Shell session has ViewEntireForest enabled. Enter **Get-ADServerSettings** to view the current Active Directory Server settings. Enter **Set-ADServerSettings-ViewEntireForest \$true** to set the recipient scope to the entire forest.

You'll find a comprehensive discussion of Exchange Management Shell in *Exchange Server 2016 & Exchange Online: Essentials for Administration* (Stanek & Associates, 2016). See "Working with Exchange Management Shell" in Chapter 1 and "Using Windows PowerShell with Exchange Online" in Chapter 2. *Essentials for Administration* also has details on troubleshooting and advanced techniques for connecting to Exchange 2016 without installing the management tools or using Exchange Management Shell. See Chapter 11 "Customizing & Troubleshooting the Exchange Shell."

Whenever you remotely manage Exchange services using Powershell, you are relying on the Windows PowerShell remoting features. These features are supported by the WS-Management protocol and the Windows Remote Management (WinRM) service that implements WS-Management in Windows.

Windows Management Framework includes Windows PowerShell and WinRM. Computers running Windows 8 and later, as well as Windows Server 2012 and later, include Windows Management Framework. One way to verify the availability of WinRM services and configure Windows PowerShell for remoting is to follow these steps:

1. Type **PowerShell I** in the Search box. Next, right-click the Windows PowerShell shortcut in the search results and select Run As Administrator.
2. The WinRM service is configured for manual startup by default. You must change the startup type to Automatic and start the service on each computer you want to work with. At the PowerShell prompt, you can verify that the WinRM service is running by using the following command:

```
get-service winrm
```

As shown in the following example, the value of the Status property in the output should be Running:

```
Status  Name          DisplayName
-----  ----          -
```

If the service is stopped, enter the following command to start the service and configure it to start automatically in the future:

```
set-service -name winrm -startuptype automatic -status running
```

3. To configure Windows PowerShell for remoting, type the following command:

```
Enable-PSRemoting -force
```

Exchange 2016 is designed to be remotely managed from domain-joined computers. If your computer is connected to a public network, you need to disconnect from the public network, connect to a domain, and then repeat this step. If one or more of your computer's connections has the Public connection type, but you are actually connected to a domain network, you need to change the network connection type in Network And Sharing Center and then repeat this step.

In many cases, you will be able to work with remote computers in other domains. However, if the remote computer is not in a trusted domain, the remote computer might not be able to authenticate your credentials. To enable authentication, you need to add the remote computer to the list of trusted hosts for the local computer in WinRM. To do so, type the following:

```
winrm s winrm/config/client '@{TrustedHosts="RemoteComputer"}'
```

where *RemoteComputer* is the name of the remote computer, such as:

```
winrm s winrm/config/client '@{TrustedHosts="MailServer12"}'
```

If you cannot connect to a remote host, verify that the service on the remote host is running and is accepting requests by running the following command on the remote host:

```
winrm quickconfig
```

This command analyzes and configures the WinRM service. If the WinRM service is set up correctly, you'll see output similar to the following:

```
WinRM already is set up to receive requests on this machine.
```

```
WinRM already is set up for remote management on this machine
```

If the WinRM service is not set up correctly, you'll see errors and need to respond affirmatively to several prompts that allow you to automatically configure remote management. When this process completes, WinRM should be set up correctly.

Whenever you use Windows PowerShell remoting features, you must start Windows PowerShell as an administrator by right-clicking the Windows PowerShell shortcut and selecting Run As Administrator. When starting Windows PowerShell from another program, such as the command prompt (cmd.exe), you must start that program as an administrator.

Chapter 14. Preparing for Exchange 2016

Exchange organizations can be deployed in several configurations, including on-premises, online and hybrid implementations. With an on-premises implementation you deploy Exchange server hardware on your network and manage all aspects of the implementation, including server configuration, organization configuration, and recipient configuration. Administrators manage Exchange using Exchange Admin Center and Exchange Management Shell. Users access Exchange using Outlook Web App and a URL provided by your organization or with Microsoft Outlook.

With an online implementation you rely on hardware and services provided by Microsoft. Here, you subscribe to Exchange Online, manage service-level settings using Office 365 Admin Center, and manage the organization and recipient configuration using Exchange Admin Center. Users access Exchange using Outlook Web App and a URL provided by Microsoft or with Microsoft Outlook.

With a hybrid implementation you integrate on-premises and online components. Here, the on-premises and Exchange Online organizations have a shared domain namespace, and mail is securely routed between them. These organizations share a unified global address list, free/busy data, and calendar data. Administrators manage Exchange using a combination of the on-premises and online tools. Users can access Exchange using Outlook Web App and the same URL whether their mailbox is stored on premises or online. Users also can access Exchange using Microsoft Outlook.

When you use an online implementation, Microsoft manages the hardware configuration and ensures availability. Otherwise, you are responsible for any on-premises hardware. Before you deploy an on-premises or hybrid implementation of Exchange 2016, you should carefully plan the messaging architecture. Every Exchange implementation has three layers in its architecture network, directory and messaging.

The network layer provides the foundation for computer-to-computer communications and essential name resolution features. The network layer has both physical and logical components. The physical components include the IP addresses, the IP subnets, local area network (LAN) or wide area network (WAN) links used by messaging systems as well as the routers that connect these links, and firewalls that protect the infrastructure. The logical components are the Domain Name System (DNS) zones that define the naming boundaries and contain the essential resource records required for name resolution.

The directory layer provides the foundation necessary for authentication, authorization, and replication. The directory layer is built on the Active Directory directory service and has both physical and logical components. The physical components include the domain controllers, Global Catalog servers, and site links used for authentication, authorization, and replication. The logical components include the Active Directory forests, sites, domains, and organizational units that are used to group objects for resource sharing, centralized management, and replication control. The logical

components also include the users and groups that are part of the Active Directory infrastructure.

The messaging layer provides the foundation for messaging and collaboration. The messaging layer has both physical and logical components. The physical components include individual Exchange servers that determine how messages are delivered and mail connectors that determine how messages are routed outside an Exchange server's routing boundaries. The logical components specify the organizational boundaries for messaging, mailboxes used for storing messages, public folders used for storing data, and distribution lists used for distributing messages to multiple recipients.

Whether you are deploying Exchange Server for the first time in your organization or upgrading to Exchange Server 2016 from an earlier release of Exchange Server, you need to closely review each layer of this architecture and plan for required changes. As part of your implementation planning, you also need to look closely at the roles your Exchange servers will perform and modify the hardware accordingly to meet the requirements of these roles on a per-server basis. Exchange Server is a complex messaging platform with many components that work together to provide a comprehensive solution for routing, delivering, and accessing email messages, voice-mail messages, faxes, contacts, and calendar information.

Designing the Exchange Server Organization

With Exchange Server Setup, you can deploy servers with the Mailbox or Edge Transport role. Prior to setup and configuration, you need to decide how you will use Exchange Server 2016, what roles you will deploy, and where you will locate those roles. Afterward, you can plan for your deployment and then roll out Exchange Server 2016.

As part of your planning and testing, you can use the Exchange Server Deployment Assistant and the Exchange Remote Connectivity Analyzer. Both are web-based tools that provide step by step guidance. The Deployment Assistant, which can help you plan on-line, on-premises and hybrid deployments, is available at <http://go.microsoft.com/fwlink/p/?LinkId=277105> and the Connectivity Analyzer, which can help you diagnose connectivity issues, is available at <https://testexchangeconnectivity.com>.

On-premises implementations of Exchange Server have three layers in their architecture: a network layer, directory layer, and messaging layer. The messaging layer is where you define and deploy the Exchange Server roles. The Exchange servers at the core of the messaging layer can operate in the following roles:

- **Mailbox Server** A primary mail server that hosts mailboxes, public folders, and related messaging data, such as address lists, resource scheduling, and meeting items. Mailbox servers accept connections to Exchange Server from a variety of clients and hosts the protocols used by all clients when checking messages. On the local network, Outlook MAPI clients are connected directly to the Mailbox server to check mail using SMTP. Remote users can check their mail over the Internet by using Outlook Anywhere, Outlook Web App, Exchange ActiveSync, POP3, or IMAP4.
- **Edge TransportServer** An additional mail routing server that routes mail into and out of the Exchange organization. This server is designed to be deployed in an organization's perimeter network and is used to establish a secure boundary between the organization and the Internet. This server accepts mail coming into the organization from the Internet and from trusted servers in external organizations, processes the mail to protect against some types of spam messages and viruses, and routes all accepted messages to a Mailbox server inside the organization.

Exchange 2016 Mailbox servers also host:

- **Client Access services** Middle-tier services that accept connections to Exchange Server from a variety of clients. These services host the protocols used by all clients when checking messages. On the local network, Outlook MAPI clients are connected directly to the server to check mail using SMTP. Remote users can check their mail over the Internet by using MAPI over HTTP, RPC over HTTP, Outlook Web App, Exchange ActiveSync, POP3, or IMAP4.
- **Unified Messaging services** Middle-tier services that integrate a private branch

exchange (PBX) system with Exchange Server 2016, allowing voice messages and faxes to be stored with email in a user’s mailbox. Unified messaging supports call answering with automated greetings and message recording, fax receiving, and dial-in access. With dial-in access, users can use Outlook Voice Access to check voice mail, email, and calendar information; to review or dial contacts; and to configure preferences and personal options. To receive faxes, you need an integrated solution from a Microsoft partner.

- **Transport services** Mail routing services that handle mail flow, routing, and delivery within the Exchange organization. These services process all mail that is sent inside the organization before it is delivered to a mailbox in the organization or routed to users outside the organization. Processing ensures that senders and recipients are resolved and filtered as appropriate, content is filtered and has its format converted if necessary, and attachments are screened. To meet any regulatory or organizational compliance requirements, the Mailbox server can also record, or journal, messages and add disclaimers to them.

The Mailbox and Edge Transport roles are the building blocks of on-premises Exchange organizations. Table 14-1 provides an overview of the basic processor configurations I recommend for these roles. Processors can have multiple cores. Following the configurations shown in the table, I recommend that you build Edge Transport servers for scaling out and Mailbox servers for scaling up.

TABLE 14-1 Recommended Configurations for Exchange Server Roles

SERVER ROLE	MINIMUM PROCESSORS	RECOMMENDED PROCESSORS	BUILD FOR
Edge Transport	1	2-4	Scale out
Mailbox	1-2	2-8	Scale up

One of the most basic Exchange organizations you can create is one that includes a single Exchange server that provides the Mailbox Server role. A server with this role is the minimum required for routing and delivering messages to both local and remote messaging clients. For added security and protection, you can deploy the Edge Transport server role in a perimeter network on one or more separate servers.

Although a basic implementation of Exchange Server might include only one server, you’ll likely find investing in multiple servers is more effective in terms of time, money, and resources. Why? High availability is integrated into the core architecture of Exchange Server 2016 and can be easily enabled.

With the Mailbox Server role, you can configure automatic failover by making the Mailbox servers members of the same database availability group. Each Mailbox server in the group can then have a copy of the mailbox databases from the other Mailbox servers in the group. Each mailbox database can have up to 16 copies, and this means you can have up to 16 Mailbox servers in a database availability group as well.

Planning for High Availability

The underlying functionality of a Mailbox server is similar to that of a database server. Every mailbox-enabled recipient defined in the organization has a mailbox that is used to store messaging data. Groups of related mailboxes are organized using databases, and each database can have one or more database copies associated with it.

With early releases of Exchange, you needed dedicated hardware for clustered Mailbox servers, those servers could not run other roles, and failover occurred at the server level. Microsoft engineered Exchange 2016 to provide continuous availability without the need for dedicated hardware for clustering. Microsoft did this by:

- [Integrating key components of Windows clustering into Exchange Server and adding components that allow Exchange Server to automatically manage the clustering functions.](#)
- [Integrating key features of Cluster Continuous Replication \(CCR\) and Standby Continuous Replication \(SCR\) into Exchange Server and making related features available through database availability groups.](#)

Using the built-in clustering features, you can create a fully redundant Exchange organization using only two Mailbox servers, with each server configured as part of the same database availability group. You would also need a witness server for the database availability group, which doesn't have to be an Exchange server.

The underlying technology built into database availability groups is the key ingredient that makes high availability possible. The related framework ensures failover clustering occurs in the background and doesn't normally require administrator intervention. As a result, Exchange Server 2016 doesn't need, or use, a cluster resource dynamic-link library (DLL) and uses only a small portion of the Windows clustering components, including heartbeat capabilities and a cluster database.

Database availability groups use continuous replication to achieve high availability. With continuous replication, Exchange Server 2016 uses its built-in asynchronous replication technology to create copies of mailbox databases and then keeps the copies up to date using transaction log shipping and replay. Lagged copies can automatically play down log files to automatically recover from certain types of issues. For example, if Exchange detects that a low disk space threshold has been reached, Exchange automatically replays the logs into the lagged copy to play down the log files. If Exchange detects that page patching is required, Exchange automatically replays the logs into the lagged copy to perform page patching. If Exchange detects that there are fewer than three available healthy copies (whether active or passive) for more than 24 hours, Exchange automatically replays the logs into the lagged copy to play down the log files.

Any server in a group can host a copy of a mailbox database from any other server in the group. When a server is added to a group, it works with other servers in the group to provide automatic recovery from failures that affect mailbox databases, including server

failure, database corruption, disk failure, and network connectivity failure. Although Exchange 2010 used a scheduled script to alert you that only a single copy of a database was available, this functionality is now integrated into the core architecture along with other managed availability features for internal monitoring and recovery.

When you create a database availability group, Exchange adds an object to Active Directory representing the group. This object stores information about the group, including details about servers that are members of the group. When you add the first server to the group, a failover cluster is created automatically and the heartbeat is initiated. As you add member servers to the group, the heartbeat components and the cluster database are used to track and manage information about the group and its member servers, including server status, database mount status, replication status, and mount location.

Placement and sizing of availability groups is something you should also consider. Focus on whether you need to use a site resilient plan with a datacenter pair as part of your design. Large enterprises that want to achieve high availability and site resiliency will want to deploy an availability group in each of at least two datacenters that are well-connected (meaning you have low, round-trip latency across the network connecting the datacenters).

Generally, you'll want to deploy a single availability group that stretches across the datacenters and scale the number of servers in the group as appropriate. For example, with two datacenters, you deploy two servers in each datacenter and make each server a member of the same availability group. This availability group then has four member servers and you can add servers to each location as the needs of your organization grows.

Although Active Directory sites can stretch across multiple datacenters, you'll usually want each datacenter to have its own Active Directory site. Having servers in different Active Directory sites is required for transport site resilience anyway. Why? Exchange 2016 achieves site resilience using Shadow Redundancy and Safety Net, which both require more than one Active Directory site to achieve transport site resilience. Shadow redundancy ensures that messages are protected from loss the entire time they are in transit by creating a copy of a message and retaining this copy while a message is in transit. Safety Net maintains a queue of messages that were recently delivered to recipients.

Planning Exchange Databases and Storage

Because Exchange databases are represented at the organization level, they are effectively disconnected from the servers on which they are stored, which makes it easier to move databases from one server to another. However, it also means you can work with databases in ways that were not possible with early releases of Exchange and that there are also several requirements when working with databases. Keep the following in mind when working with databases in Exchange Server 2016:

- Database names must be unique throughout your Exchange organization. This means you cannot name two databases identically even if they are on two different Mailbox servers.
- Every mailbox database, except copies, have a different globally unique identifier (GUID). Copies of a database have the same GUID as the original database.
- Mailbox servers that are part of the same database availability group do not require cluster-managed shared storage. However, the full paths for all database copies must be identical on host Mailbox servers.

NOTE Exchange no longer supports public folder databases. Beginning with Exchange 2013, public folder data was moved into specially configured mailboxes that the public folder hierarchy and content. Like traditional mailboxes, these special mailboxes for public folders are stored in mailbox databases and are replicated as part of any database availability group you configure.

For a successful deployment of a Mailbox server, the storage subsystem must meet the storage capacity requirements and must be able to perform the expected number of input/output (I/O) operations per second. Storage capacity requirements are determined by the number of mailboxes hosted on a server and the total storage size allowed per mailbox. For example, if a server hosts 2,500 mailboxes that you allow to store up to 2 gigabytes (GB) each, you need to ensure the storage system capacity can be scaled up to support this.

I/O performance of the storage subsystem is measured in relation to the latency (delay) for each read/write operation to be performed. The more mailboxes you store on a specific drive or drive array, the more read/write operations there are performed and the greater the potential delay. To improve performance, you can use multiple mailbox databases on separate disks. You might also want to store databases with their transaction log files on separate disk drives, such that database A and related logs are on disk 1, database B and related logs are on disk 2, and so on. In some scenarios, you might want the databases and logs to be on separate disks.

Exchange 2016 supports either NTFS or ReFS for partitions containing databases, content indexes and transaction logs. The on-disk storage engine for ReFS is very different from the on-disk storage engine for NTFS. ReFS uses B+ tree structures to represent all information on a disk. B+ tree structures scale well and simplify the

architecture. ReFS has many other enhancements over NTFS, including improved reliability and failure recovery. Although ReFS is designed to be used with Storage Spaces, you don't have to use Storage Spaces to gain many of the benefits of this enhanced architecture.

I/O performance in Exchange Server 2016 running on 64-bit architecture is improved substantially over 32-bit architecture. On Mailbox servers, a 64-bit architecture enables a database cache size of up to approximately 90 percent of total random access memory (RAM). A larger cache increases the probability that data requested by a client will be serviced out of memory instead of by the storage subsystem.

Because Exchange 2016 allows a server to host multiple databases on the same volume, you don't need separate volumes for each database copy. As part of your planning, look closely at the input/output per second (IOPS) capabilities of your storage architecture and place database copies appropriately. Because active copies will use more IOPS than passive copies, you'll typically want no more than one active database copy on a volume while allowing multiple passive copies. For example, if you're configuring a four-server database availability group, you could configure storage so that each server has a primary storage volume for the active database copy and passive copies of the databases from the other servers.

As part of the preferred architecture, Microsoft recommends that active database copies are distributed equally across the availability group. This configuration helps to ensure that the workload is distributed across all servers and that the full stack of client connectivity, transport and replication services are being validated during normal operations.

Exchange 2016 is optimized so that servers can use large disks with 2 to 8 terabytes of storage efficiently. With very large volumes like these, you'll want to consider using ReFS rather than NTFS. ReFS uses 128-bit file identifiers which support a larger number of files and directories, and hierarchical allocators which can optimally allocate storage more quickly than NTFS.

If you use ReFS with Exchange Server, you'll also want to:

- Ensure the integrity feature is disabled. Use the `Get-FileIntegrity` cmdlet to check the status of the integrity feature and the `Set-FileIntegrity` cmdlet to change the status of the integrity feature.
- Ensure Mailbox servers configured in availability groups are configured so that the auto reseed feature formats disks with ReFS. Use the command `Set-DatabaseAvailabilityGroup DagName -FileSystem ReFS`, where *DagName* is the name of the availability group.

As part of your planning, you also need to understand how Exchange 2016 uses automatic reseed to recover from disk failure, database corruption events, and other issues that require a reseed of a database copy. With automatic reseed, Exchange can automatically restore database redundancy using spare disks that have been pre-provisioned.

The larger the database, the longer it takes Exchange to reseed it. If a database is too large, it can't be reseeded in a reasonable amount of time. With a typical reseed rate of 20 MB per second, it would take Exchange:

- About 28 hours to reseed a 2-terabyte database.
- About 42 hours to reseed a 3-terabyte database.
- About 56 hours to reseed a 4-terabyte database.

@techjob

Because of this, the total reseed time may be the most important limiting factor for sizing databases. When configuring disks, at least one disk should be reserved as a hot spare to allow auto reseed to restore database redundancy. In the event of a disk failure, auto reseed activates the hot spare and initiates a database copy to reseed the database.

Planning for Client Access

As part of the architecture changes for Exchange 2016, Mailbox servers handle all of the client-related messaging tasks in an Exchange implementation. This means that not only do Mailbox servers perform all mail processing and content conversion, but they also perform authentication, proxying and redirection of client connections as appropriate.

Clients don't connect directly to back-end services on Mailbox servers. In a basic configuration, clients connect to Client Access services running on a Mailbox server and then are routed via local or remote proxy to the back-end endpoint on the Mailbox server that hosts the active copy of the mailbox database storing the user's mailbox. More typically, clients connect to your Exchange organization using the load-balanced virtual IP address of the protocol being used, which involves a series of steps:

1. After the client resolves the namespace to a load-balanced virtual IP address, the load balancer assigns the session to a Mailbox server in the load-balanced server pool.
2. Client Access services running on the Mailbox server authenticate the request and perform a service discovery by accessing Active Directory to identify the Mailbox version and location details.
3. After Client Access services running on the Mailbox server locate the user's mailbox and the active copy of the associated mailbox database, the server either proxies the request or redirects the request to the appropriate Mailbox server within the same forest.

When you are configuring Exchange 2016, you need to determine the client protocols that you want to implement. Exchange Server 2016 allows local access using Microsoft Outlook with Simple Mail Transfer Protocol (SMTP) and remote access using MAPI over HTTP, Outlook Anywhere (RPC over HTTP), Outlook Web App, and Exchange ActiveSync.

Internet Message Access Protocol 4 (IMAP4) and Post Office Protocol 3 (POP3) are available as alternatives to standard protocols. IMAP4 is a protocol for reading mail and accessing public and private folders on remote servers. POP3 is a protocol for retrieving mail from remote servers. Client Access servers provide access to free/busy data by using the Availability service, and they enable clients to download automatic configuration settings from the Autodiscover service.

Each client protocol that you want Exchange to support requires a namespace configuration. For example:

- For autodiscover: autodiscover.tvpress.com
- For HTTP clients: mail.tvpress.com
- For IMAP clients: imap.tvpress.com
- For POP3 clients: pop3.tvpress.com
- For SMTP clients: smtp.tvpress.com

Thus, you will need to configure each of these namespaces as appropriate for your organization. When you are operating Exchange out of two or more datacenters, namespaces can be said to be either:

- **Bound** With bound namespaces, users are associated with a specific datacenter and there typically are multiple namespaces: a primary and a failback namespace which generally correspond to the primary datacenter and a secondary datacenter used only during failure events. Here, when the organization has a datacenter pair, there typically are separate database availability groups operating in each datacenter with each group containing a set of mailboxes for that datacenter.
- **Unbound** With unbound namespaces, users are not associated with a specific datacenter and there is a single namespace, allowing user requests to be serviced out of any available datacenter. Here, when the organization has a datacenter pair, there typically is a single database availability group with member servers operating in each datacenter.

With unbound namespaces, you deploy a unified namespace and the virtual IP address assigned to the namespace is load balanced to distribute the workload between datacenters. In contrast, with a bound namespace, you deploy a dedicated namespace for each datacenter. Generally, unbound namespaces are preferred to bound namespaces as any available datacenter can service user requests. If you have multiple datacenters and use an unbound model, you can load balance across the datacenters using a layer 7 configuration without session affinity. Session affinity isn't required because sessions are maintained directly with the Mailbox server hosting the active copy of the mailbox database storing the user's mailbox.

Exchange 2016 uses the Active Directory infrastructure to determine its site membership and the site membership of other servers. The Microsoft Exchange Active Directory Topology service running on an Exchange server is responsible for updating the site attribute of an Exchange server in the directory.

Once a server determines its site membership, the server identifies which domain controllers and global catalogs to use for processing Active Directory queries. Because this information is available in the directory, Exchange servers don't need to use DNS to resolve a server address to a subnet associated with an Active Directory site.

Exchange 2016 Mailbox servers interact directly with Outlook clients and Active Directory. Mailbox servers use Lightweight Directory Access Protocol (LDAP) to obtain recipient, server, and organization configuration information from Active Directory. Mailbox servers accept client connections over the local network and over the Internet, returning data to clients as appropriate, including online address book files, free/busy data, calendar schedules, and client profile settings.

Some clients use POP3 or IMAP4 connections to communicate with the Exchange server. Other clients use SMTP, POP3, or IMAP4 to communicate with the Exchange server. The IIS server host running on each Mailbox server manages the web applications for Outlook Web App, Exchange Active Sync, Exchange Admin Center, and

PowerShell.

Outlook clients on the corporate network access the Mailbox server to send and retrieve messages. Outlook clients outside the corporate network can use Outlook Anywhere (RPC over HTTP) or MAPI over HTTP to access Mailbox servers. Regardless of whether they are on or outside the corporate network, Outlook clients access public folder data using either RPC over HTTP or MAPI over HTTP. To retrieve a user's Active Directory information, Mailbox servers use LDAP or Name Service Provider Interface (NSPI). By default, communications with domain controllers and global catalogs are encrypted.

NOTE RPC connections are made directly to the MAPI RPC connection point on the Mailbox server and the NSPI endpoint on the Mailbox server. For directory information, Outlook communicates with an NSPI endpoint located on the Mailbox server. NSPI communicates with the Active Directory driver, which then communicates with Active Directory.

It's important to point out that Microsoft recommends using a split-brain approach to DNS infrastructure with Exchange. Split-brain DNS enables different IP addresses to be returned for a particular namespace depending on where the clients is located. If the client is on the internal network, internal IP addresses are used. If the client is outside the internal network, external IP addresses are used. Using split-brain DNS also simplifies the configuration of Exchange virtual directories, as you can use the same values for internal and external URLs. If you don't use split-brain DNS, you will need to specify different internal and external URLs.

Planning to Support Transport Services

The Transport services on Mailbox servers and the Edge Transport role perform similar tasks. You use both for messaging routing, and both have a similar set of filters to protect an organization from spam and viruses. The key difference is in where you place servers with these roles. You place a Mailbox server in the internal network and configure it as a member of the organizational domain. If you use a server with the Edge Transport role, you place it in the organization's perimeter network, and you do not configure it as a member of the organizational domain.

Mailbox servers and Edge Transport servers cannot have the SMTP or Network News Transfer Protocol (NNTP) service installed separately. Although you install Edge Transport servers outside the Active Directory forest, the server must have a DNS suffix configured, and you must be able to perform name resolution from the Edge Transport server to any Mailbox servers.

TIP Transports store all incoming mail in a database file called mail.que until the transport verifies that all of the next hops for that message have been completed. This database has an associated transaction log in which changes are first committed. If you are using an Exchange Server's internal drive(s) for storage in a high-volume environment in which one million or more messages are persisted, you should consider placing the database and the transaction log on separate disks for optimal performance and fault tolerance. With Storage Area Networks (SANs), it might not be immediately apparent whether disks are physically separate. This is because the volumes you see are logical references to a portion of the storage subsystem. In this case, you might be able to use the Storage Manager For SANs console or a similar tool to help you select logical unit numbers (LUNs) that are on physically separate disks.

MORE INFO Transports have many different queues for messages. These queues are all stored in a single Extensible Storage Engine (ESE) database called mail.que. By default, this database is located in %ExchangeInstallPath%\TransportRoles\data\Queue. Thanks to shadow redundancy, the deletion of a message in the database is delayed until the transport verifies that all of the next hops for that message have completed delivery. If any of the next hops fail before reporting back successful delivery, the message is resubmitted for delivery to that next hop.

Both Mailbox servers and Edge Transport servers can perform protocol logging and message tracking. Only Mailbox servers perform content conversion to format messages for recipients. Protocol logging allows you to verify whether a protocol is performing as expected and whether any issues need attention. Because this feature is designed for troubleshooting, it is disabled by default. Message tracking creates logs that track messages sent and received. Incoming mail from the Internet is converted to Summary Transport Neutral Encoding Format (STNEF) prior to being delivered. STNEF

messages are always MIME-encoded and always have a Content-Transfer-Encoding value of Binary. Because content conversion is performed in the temp folder, you can improve performance by ensuring that the temp folder is not on the same physical disk as the paging file and operating system.

The transport pipeline used by Exchange 2016 is different from the transport pipeline for Exchange 2010 and has the following key components:

- [Front End Transport service](#)
- [Transport service](#)
- [Mailbox Transport Submission service](#)
- [Mailbox Transport Delivery service](#)

The Front End Transport service is a client access service that proxies all inbound and outbound SMTP traffic for Exchange 2016. Mailbox servers also have a separate Transport service, which handles back-end tasks for SMTP mail flow as well as message categorization and message content inspection. The Transport service doesn't communicate directly with mailbox databases. Instead, the Mailbox Transport Submission and Mailbox Transport Delivery services are used to provide separate mail submission and delivery processes.

The basic submission process works like this:

1. The Mailbox Transport Submission service receives SMTP messages from the Transport service on the local Mailbox server or on other Mailbox servers
2. The Mailbox Transport Submission service connects to the local mailbox database.
3. The Mailbox Transport Submission service uses RPC to deliver the message.

The basic delivery process works like this:

1. The Mailbox Transport Delivery service connects to the local mailbox database using RPC to retrieve messages.
2. The Mailbox Transport Delivery service submits messages over SMTP to the Transport service on the local Mailbox server or on other Mailbox servers.
3. The Transport service routes messages using SMTP.

Messages from inside the organization enter the transport pipeline through a Receive connector, from the Mailbox Transport Delivery service, from the Pickup or Replay directories, or from agent submission. Messages from outside the organization enter the transport pipeline through a Receive connector in the Front End Transport service and are then routed to the Transport service.

Planning for Unified Messaging

Unified messaging allows you to integrate voice mail, fax, and email functionality so that the related data can be stored in a user's Exchange mailbox. To implement unified messaging, your organization must have a PBX that is connected to the LAN, and you must deploy Mailbox servers running Exchange Server 2016. After deployment, the Unified Messaging service running on a Mailbox server has the job of providing call answering, fax receiving, subscriber access, and auto-attendant features that allow access to content over the telephone and storage of content received from the PBX. However, it is the job of the Unified Messaging Call Router service running on Mailbox servers to provide call routing and proxy services that allow calls to be connected.

Although some current PBXs, referred to as *IP-PBXs*, are Internet Protocol-capable, all other PBXs require a separate Internet Protocol/Voice over Internet Protocol (IP/VoIP) gateway to connect to the LAN. After you connect a PBX to the LAN, you can link it to Exchange by deploying and appropriately configuring the Unified Messaging service. The Desktop Experience feature, which is required to install Exchange server, provides the Microsoft Speech service, Microsoft Windows Media Encoder, and Microsoft Windows Media Audio Voice Code components used by the Unified Messaging service.

The Unified Messaging service doesn't perform a great deal of I/O operations, and the primary potential bottlenecks for this service are the processors, memory, and network. Disk I/O operations for this service are primarily limited to accessing routing details and dial plans, which include auto-attendant and mail policy settings.

If you are planning to use Unified Messaging in a hybrid Exchange implementation, you'll also need to configure session board controllers (SBCs). SBCs have two IP interfaces: one for your network and another that connects over the Internet. Your VoIP, IP-PBX, and SBC components must be configured to communicate with your Mailbox servers. You also must create and configure a Unified Messaging IP gateway to represent each deployed device.

Integrating Exchange with Active Directory

Exchange 2016 makes extensive use of Active Directory. Each Exchange server must access Active Directory to retrieve information about recipients and the Exchange organization. Exchange 2016 only works with read-writeable domain controllers.

How Mailbox Servers use Active Directory

Mailbox servers are service locations for email messages, voice-mail messages, and faxes. For outgoing mail, Mailbox servers can access Active Directory to retrieve information about the location of Mailbox servers in their site. Then they can use this information to forward messages for routing.

The Transport service running on Mailbox servers contacts Active Directory for message categorization. The Categorizer queries Active Directory to perform recipient lookup, retrieves the information needed to locate a recipient's mailbox (according to the mailbox store in which it is created), and determines any restrictions or permissions that might apply to the recipient. The Categorizer also queries Active Directory to expand the membership of distribution lists and to perform the LDAP query processing when mail is sent to a dynamic distribution list.

After the Categorizer determines the location of a mailbox, the Transport service uses Active Directory site configuration information to determine the routing topology and locate the site of the mailbox. If the mailbox is in the same Active Directory site as the Mailbox server, the Transport service delivers the message directly to the user's mailbox. If the mailbox is in a different Active Directory site from the Mailbox server, the Transport service delivers the message to a Mailbox server in the remote Active Directory site.

Mailbox servers store all configuration information in Active Directory. This configuration information includes the details of any transport or journaling rules and connectors. When this information is needed, a Mailbox server accesses it in Active Directory.

Mailbox servers also store configuration information about mailbox users, mailbox stores, agents, address lists, and policies in Active Directory. Mailbox servers retrieve this information to enforce recipient policies, mailbox policies, system policies, and global settings.

Client Access services running on Mailbox servers receive connections from local and remote clients. At a high level, when a user connection is received, the Client Access services contacts Active Directory to authenticate the user and to determine the location of the user's mailbox. If the user's mailbox is in the same Active Directory site as the originating Mailbox server, the user is connected to the mailbox. If the user's mailbox is in an Active Directory site other than the one the originating Mailbox server is located in, the connection is redirected to a Mailbox server in the same Active Directory site as

the user's mailbox.

At least one of your Mailbox servers in each site must be designated as Internet-facing. The Internet-facing server proxies requests from Outlook Web App, Exchange ActiveSync, and Exchange Web Services to the Mailbox server closest to the user's mailbox.

When deployed, the Unified Messaging service running on Mailbox servers accesses Active Directory to retrieve global configuration information, such as dial plans and IP gateway details. When a message is received by the Unified Messaging service, the service searches for Active Directory recipients to match the telephone number to a recipient address. When the service has resolved this information, it can determine the location of the recipient's mailbox and then submit the message to the appropriate Mailbox server for submission to the mailbox.

How Edge Transports use Active Directory

You deploy Edge Transport servers in perimeter networks to isolate them from the internal network. As such, they are not members of the internal domain and do not have direct access to the organization's internal Active Directory servers for the purposes of recipient lookup or categorization. Thus, unlike the Transport service on Mailbox servers, Edge Transport servers cannot contact an Active Directory server to help route messages.

To route messages into the organization, an administrator can configure a subscription from the Edge Transport server to the Active Directory site that allows it to store recipient and configuration information about the Exchange organization in its AD LDS data store. After an Edge Transport server is subscribed to an Active Directory site, it is associated with the Mailbox servers in that site for the purpose of message routing. Thereafter, Mailbox servers in the organization route messages being delivered to the Internet to the site associated with the Edge Transport server, and Mailbox servers in this site relay the messages to the Edge Transport server. The Edge Transport server, in turn, routes the messages to the Internet.

The EdgeSync service running on Mailbox servers is a one-way synchronization process that pushes information from Active Directory to the Edge Transport server.

Periodically, the EdgeSync service synchronizes the data to keep the Edge Transport server's data store up to date. The EdgeSync service also establishes the connectors needed to send and receive information that is being moved between the organization and the Edge Transport server and between the Edge Transport server and the Internet. The key data pushed to the Edge Transport server includes:

- [Accepted and remote domains](#)
- [Valid recipients](#)
- [Safe senders](#)
- [Send connectors](#)
- [Available Mailbox servers](#)

- Available SMTP servers
- Message classifications
- TLS Send and Receive Domain Secure lists

After the initial replication is performed, the EdgeSync service synchronizes the data periodically. Configuration information is synced once every hour, and it can take up to one hour for configuration changes to be replicated. Recipient information is synced once every four hours, and it can take up to four hours for changes to be replicated. If necessary, administrators can initiate an immediate synchronization using the Start-EdgeSynchronization cmdlet in Exchange Management Shell.

NOTE During synchronization, objects can be added to, deleted from, or modified in the Edge Transport server's AD LDS data store. To protect the integrity and security of the organization, no information is ever pushed from the Edge Transport server's AD LDS data store to Active Directory.

Integrating Exchange 2016 Into Existing Organizations

Existing Exchange Server 2010 and Exchange Server 2013 installations can coexist with Exchange Server 2016 installations. Generally, you do this by integrating Exchange Server 2016 into your existing Exchange Server 2010 or Exchange Server 2013 organization. Integration requires the following:

- [Preparing Active Directory and the domain for the Active Directory changes that will occur when you install Exchange Server 2016.](#)
- [Configuring Exchange Server 2016 so that it can communicate with servers running Exchange Server 2010 and Exchange Server 2013.](#)
- [Ensuring any installations of Exchange 2010 are running Service Pack 3 with RU11 at a minimum and any installations of Exchange 2013 are running Cumulative Update 10 or later.](#)

You cannot upgrade existing Exchange 2010 or Exchange 2013 servers to Exchange 2016. You must install Exchange Server 2016 on new hardware, and then move the mailboxes from your existing installations to the new installation. See the “Moving to Exchange Server 2016” section later in this chapter for more details.

As an alternative to coexistence, you can deploy a new Exchange 2016 organization. After you deploy a new Exchange 2016 organization, you can't add servers that are running earlier versions of Exchange to the organization. Adding earlier versions of Exchange to a new Exchange 2016 organization is not supported.

Coexistence and Active Directory

As Exchange Server 2016 contains schema changes and other Active Directory updates, you might want to prepare Active Directory and the domain for these changes prior to installing Exchange 2016 for the first time, especially in a large enterprise.

To do this, follow these steps:

1. [Prepare the schema by running the following command prior to executing the Exchange 2016 Setup:](#)

setup.exe /PrepareSchema /IAcceptExchangeServerLicenseTerms

This command connects to the schema master and imports the LDAP data interchange format files that are used to update the schema with Exchange 2016 specific attributes. Optionally, use the `/DomainController` parameter to specify the name of the schema master. You must run this command on a 64-bit computer in the same domain and site as the schema master. If schema needs to be updated and you haven't previously prepared schema, you must ensure the account you use is delegated membership in the Schema Admins group. Wait for the changes to replicate before continuing.

2. Prepare Active Directory for Exchange 2016 by running the following command prior to executing the Exchange 2016 Setup:

setup.exe /PrepareAD /IAcceptExchangeServerLicenseTerms

You must run this command in the same domain and site as the schema master. This computer must be able to connect to all domains in the forest on TCP port 389. To run this command, you must be a member of the Domain Admins groups for the local domain or the Enterprise Admins group. Wait for the changes to replicate before continuing.

The PrepareAD option performs a number of tasks:

- Creates the Microsoft Exchange container and the Exchange organization container in the directory if they don't exist, such as when you are installing a new Exchange organization. Here, you must set a name for the organization using the /OrganizationName parameter.
- Verifies that the schema has been updated for Exchange 2016. It does this by checking objectVersion property for the Exchange configuration container and ensuring the value is 16210 or higher. The command also sets the Exchange product ID of the Exchange organization to that of the version you are installing. The base value for Exchange 2016 RTM is 15.01.0225.042. This value is incremented when you deploy Cumulative Updates to Exchange. Here, 15 is the major version number, 1 is the minor version number and 225.42 is the build.
- Creates any containers that are required in Active Directory for Exchange 2016, creates the default Accepted Domains entry if a default was not previously set, and imports the Rights.ldf file to add the extended rights required for Exchange to the directory.
- Creates the Microsoft Exchange Security Groups organizational unit in the root domain of the forest and then creates the following management role groups used by Exchange to this organizational unit if these haven't been previously created: Compliance Management, Delegated Setup, Discovery Management, Help Desk, Hygiene Management, Organization Management, Public Folder Management, Recipient Management, Records Management, Server Management, UM Management, and View-Only Organization Management. As necessary, also adds these groups to the otherWellKnownObjects attribute on the Exchange Services Configuration container.
- Creates the Unified Messaging Voice Originator contact in the Microsoft Exchange System Objects container of the root domain and then prepares the local domain for Exchange 2016.

3. The domain in which you ran **setup.exe /PrepareAD** is already prepared. For all other domains that will have mail-enabled users or in which you will install Exchange 2016, you must log in and run **setup.exe /PrepareDomain /IAcceptExchangeServerLicenseTerms** . You also can specify the name of the domain in which you want to run the command, such as **setup.exe /PrepareDomain:Tech.Imaginedlands.com/IAcceptExchangeServerLicenseTe**
Alternatively, you can run **setup.exe**

/PrepareAllDomains/IAcceptExchangeServerLicenseTerms to prepare all domains in the forest. To run this command, you normally must be a member of the Domain Admins groups for the local domain or the Enterprise Admins group. However, if the domain was created after running /PrepareAD, the account you use must be a member of the Exchange 2016 Organization Management role group and the Domain Admins groups in the domain.

For new organizations, this command creates the Microsoft Exchange System Objects container and sets its permissions. For all organizations, this command:

- Sets the objectVersion property in the Microsoft Exchange System Objects container so that it references the version of domain preparation for Exchange 2016, which is 16210 or higher.
- Creates a domain global group in the current domain called Exchange Install Domain Servers and adds this group in the Microsoft Exchange System Objects container as well as the Exchange Servers group in the root domain.
- Assigns permissions in the domain for the Exchange Servers group and the Organization Management group.

NOTE Want to determine the Exchange version number? Use Exchange Management Shell with either of the following commands:

```
Get-Command ExSetup.exe | % {$_.FileVersionInfo}
```

```
Get-ExchangeServer | Format-List Name, Edition, AdminDisplayVersion
```

Although Exchange Server 2016 Setup can perform these processes for you during the upgrade, the changes can take some time to replicate throughout a large organization. By performing these tasks manually, you can streamline the upgrade process. You also can ensure the tasks are run with accounts that have appropriate permissions.

To verify that schema was updated for Exchange 2016, you can use ADSI Edit to check the following properties:

- In the Schema naming context, confirm that the rangeUpper property on ms-Exchange-Schema-Version-Pt is set to 15317 or higher.
- In the Configuration naming context, confirm that the objectVersion property in the CN=*OrganizationName*,CN=Microsoft Exchange,CN=Services,CN=Configuration,DC=*DomainName* container is set 16210 or higher.
- In the Default naming context, confirm that the objectVersion property in the Microsoft Exchange System Objects container under DC=*RootDomainName* is set to 16210 or higher.

As a prerequisite for installing Exchange 2016, Active Directory must be at Windows Server 2008 forest functionality mode or higher. Additionally, the schema master for the Active Directory forest along with at least one global catalog server and at least one domain controller in each Active Directory site must be running one of the following operating systems:

- [Windows Server 2008 Standard, Enterprise or Datacenter](#)
- [Windows Server 2012 Standard or Datacenter](#)
- [Windows Server 2012 R2 Standard or Datacenter](#)
- [Windows Server 2016 Standard or Datacenter](#)

When you deploy IPv6, Exchange 2016 servers can send data to, and receive data from, devices, servers, and clients that use IPv6 addresses. However, Exchange 2016 only supports IPv6 when IPv4 is also installed and enabled. Exchange requires that IPv4 be enabled even if you don't use IPv4 addressing.

Configuring Exchange 2016 for Coexistence

When managing Exchange servers, you should use the administrative tools for that Exchange Server version. The Exchange Admin Center and Exchange Management Shell are the primary management tools for Exchange 2016. Mailboxes located on Exchange 2010 and Exchange 2013 servers are also displayed in the Exchange Admin Center.

You can manage the Exchange 2010 or 2013 mailbox properties using the Exchange Admin Center or Exchange Management Shell. You can use either tool to move mailbox recipients from Exchange 2010 or Exchange 2013 to Exchange 2016.

Beginning with Exchange 2016, MAPI over HTTP is the preferred access protocol for Outlook clients. MAPI over HTTP includes enhancements that make connections more reliable and stable, and also makes it easier for clients to recover from transport errors. Built-in pause and resume functions as well as other features allow clients to change networks and resume from sleep mode while maintaining the same server context.

Although Exchange 2016 is designed so that Outlook clients can use either Outlook Anywhere (RPC over HTTP) or MAPI over HTTP, keep the following in mind:

- [Outlook 2016 clients use MAPI over HTTP by default and can also use RPC over HTTP.](#)
- [Outlook 2013 clients must have Service Pack 1 or later to use MAPI over HTTP, and otherwise can only use RPC over HTTP.](#)
- [Outlook 2010 clients must have Service Pack 2 or later as well as KB2956191 and KB2965295 to use MAPI over HTTP, and otherwise can only use RPC over HTTP.](#)

MAPI over HTTP is enabled by default in most Exchange 2016 deployments, including new deployments, when upgrading to Exchange 2016 from Exchange 2010 and mixed environments with Exchange 2010 and Exchange 2016 servers. However, MAPI over HTTP is not enabled automatically when you have previously deployed Exchange 2013. As shown in the following example, you can use the Set-OrganizationConfig cmdlet with -MapiHttpEnabled to enable MAPI over HTTP:

```
set-organizationconfig mapihttpenabled $true
```

As with other protocols, MAPI over HTTP relies on internal and external virtual directories to be created. These virtual directories act as endpoints for client connections. You can use the Set-MapiVirtualDirectory cmdlet to specify the internal

and external URLs. You can determine whether MAPI over HTTP is working properly using `-ProbeIdentity OutlookMapiHttpSelfTestProbe` with `Test-OutlookConnectivity` as shown in the following example:

```
test-outlookconnectivity -probeidentity outlookmapihttpselftestprobe
```

MAPI over HTTP is working properly when “Succeeded” is listed as the result.

If you need to control whether a client can use MAPI over HTTP, use the `Set-CasMailbox` cmdlet with `-MapiHttpEnabled` to enable or disable the option for the associated mailbox user. The basic syntax is:

```
set-casmailbox userormailboxname -mapihttpenabled {$true | $false}
```

In the following example, you disable MAPI over HTTP for `williams@imaginedlands.com`:

```
set-casmailbox williams@imaginedlands.com -mapihttpenabled $false
```

As the mailbox setting has precedence over organization settings, the user won't be able to use MAPI over HTTP even if the protocol is enabled in the organization.

Setting the Default Offline Address Book

A new Offline Address Book (OAB) will be created when you deploy the first Exchange 2016 Mailbox server in an existing Exchange organization. All existing clients that use OAB will see this new OAB by default the next time they perform an OAB update, and they also will perform a full OAB download. If you don't want this to happen, you must configure existing mailbox databases to explicitly point to the current default OAB before you deploy the first Exchange 2016 server.

With Exchange 2010, you can do this by following these steps:

1. In Exchange Management Console, navigate to Organization Configuration, Mailbox, Database Management and then open the Mailbox Database Properties dialog box for the mailbox database you want to work with.
2. On the Client Settings tab of the Mailbox Database Properties dialog box, you'll see an entry for the Offline Address Book and a related Browse button. Use this option to explicitly set the default OAB.
3. Repeat this process for each mailbox database that you want to update.

With Exchange 2013, you can do this by following these steps:

1. In Exchange Admin Center, click Servers in the Features pane and then click Databases. Double-click the mailbox database you want to work with to open its Properties dialog box.
2. Select Client Settings. You'll see an entry for the Offline Address Book and a related Browse button. Use this option to explicitly set the default OAB.
3. Repeat this process for each mailbox database that you want to update.

You also can use Exchange Management Shell to view all mailbox databases without a default OAB explicitly set on them and then explicitly set a default OAB. Start by entering the following command:

```
Get-MailboxDatabase | Where {$_.OfflineAddressBook -eq $Null} |  
FT Name,OfflineAddressBook -AutoSize
```

If no values are returned, a default OAB is already explicitly set throughout the organization. If values are returned, you need to configure some databases with an explicitly defined default OAB. The following commands locate all mailbox databases with no default OAB defined at the database level and then set these mailbox databases to the current default OAB in the organization:

```
Get-MailboxDatabase | Where {$_.OfflineAddressBook -eq $Null} |  
Set-MailboxDatabase -OfflineAddressBook (Get-OfflineAddressBook |  
Where {$_.IsDefault -eq $True})
```

Confirm that all mailbox databases now have an explicitly defined default OAB, by re-running the first command: `Get-MailboxDatabase | Where {$_.OfflineAddressBook -eq $Null}`. The command should return no values.

Moving to Exchange Server 2016

Most organizations have existing Exchange installations. When moving Exchange 2010 or Exchange 2013 installations to Exchange Server 2016, you cannot perform an in-place upgrade. Instead, you must install new Exchange Server 2016 servers into the existing organization and then migrate to Exchange Server 2016.

Migration from Exchange 2010 or Exchange 2013 to Exchange 2016 involves installing Exchange Server 2016 on new servers and then moving the mailboxes and public folders from your existing installations to the new installation. In a migration, only mailbox and public folder data is moved, and any Exchange configuration data is not maintained.

The steps you perform to migrate from Exchange 2010 or Exchange 2013 to Exchange 2016 are as follows:

1. Plan to migrate all Exchange servers in a particular site to Exchange 2016 at the same time. You should start with Internet-accessible Active Directory sites and then migrate internal Active Directory sites. For each Exchange 2016 Mailbox server, you can configure only one Outlook Web App URL for redirection.
2. Install Exchange 2016 on new hardware and make it a member of the appropriate domain in the forest. You should install Mailbox servers first and then any Edge Transports.
3. Move Internet mail flow from Exchange 2010 or Exchange 2013 to Exchange 2016 by creating appropriate send connectors and accepted domains. You'll also need to configure default email address policy, SSL certificates and URLs for client protocols.

4. Move mailboxes and public folders from your existing Exchange 2010 or Exchange 2013 installations to the new Exchange 2016 Mailbox servers. If you move a mailbox that is part of an email address policy, the email address for the mailbox is automatically updated based on the settings in the email address policy. In this case, the new email address becomes the primary address, and the old email address becomes the secondary address.

During a migration, the version of an Exchange feature that a user sees, such as Outlook Web App, depends on where the user's mailbox is located. If the mailbox is on an Exchange 2013 server, the user sees Exchange 2013 versions of CAS features. When you move the mailbox to Exchange 2016, the user will see Exchange 2016 versions of CAS features.

REAL WORLD You move mailboxes from Exchange 2010 or Exchange 2013 to Exchange 2016 by using an online move. Perform the move from the Exchange 2016 server by using move mailbox requests, either with Exchange Management Shell or the Exchange Admin Center. You can't use the Exchange Management tools for Exchange 2010 or Exchange 2013 to move the mailboxes.

5. Once you've complete the move and have validated the configuration, you can remove unneeded Exchange 2010 or Exchange 2013 servers from the organization.

CAUTION Before removing the last Exchange 2010 or Exchange 2013 server with a particular role, you must be sure that you will never need to introduce an Exchange 2010 or Exchange 2013 server with the role again. Once you remove the last Exchange 2010 or Exchange 2013 server with a particular role, you can never add another one with that role.

Chapter 15. Deploying Exchange Server 2016

You use Exchange Server 2016 Setup to install Exchange Server roles and the Exchange management tools. You can install Exchange 2016 from media or from a download. The same media or download is used for both Exchange Server 2016 Enterprise and Exchange Server 2016 Standard.

Downloads are packaged, self-extracting, executable files. When you access the download page, click Download to start the download process. Next, copy the download to your computer for installation at a later time by clicking Save. After you copy the download to the computer on which you plan to install Exchange, you can double-click the executable file to extract the Exchange 2016 Setup components to a folder. When prompted, be sure to specify an exact folder to put all the setup components in one place. Within this folder, you'll find a program called Setup.exe. This is the Exchange Server 2016 Setup program.

You use Setup to install Exchange Server 2016 on Windows servers and to add the management tools to desktop computers. If you want to uninstall Exchange 2016, you use Programs And Features in Control Panel.

Installing New Exchange Servers

For servers deployed within the organization, you can install the Mailbox role to handle messaging transport and client access needs. While some organizations may be able to use only a single Exchange server, most organization will want to install at least two Exchange servers. Why? You can achieve high availability for the Mailbox role simply by installing two or more Mailbox servers, creating a database availability group, adding mailbox databases to this group, and then adding database copies.

You can achieve high availability for message transport simply by installing multiple Mailbox servers. Thanks to the shadow redundancy feature, a message that is submitted to a Mailbox server is stored in the transport database until the transport server verifies that all of the next hops for that message have completed delivery. If the next hop doesn't report successful delivery, the message is resubmitted for delivery. In addition, when messages are in the transport dumpster, they aren't removed until they are replicated to all the appropriate mailbox databases.

For message transport, install at least one Mailbox server for each group of Active Directory sites that are well connected on a common LAN. For example, if the organization consists of sites A and B, which are well connected on a common LAN, and sites C and D, which are well-connected on a common LAN, with wide area network (WAN) links connecting sites A and B to sites C and D, a minimal implementation would be to have Mailbox servers only in site A and site C. However, Microsoft recommends that you have Mailbox servers in each Active Directory site with mail-enabled clients.

Because you install Edge Transport servers outside the Active Directory forest, you can deploy additional Edge Transports at any time. By configuring multiple Edge Transport servers, you can ensure that if one server fails, Edge Transport services continue. If you also configure your Edge Transport servers with round-robin DNS, you can load balance between them.

REAL WORLD If you are installing Exchange Server on a new network, such as one for a new company or a development environment, be sure that you've properly configured Active Directory and DNS before installing Exchange Server. You need to create a domain. Typically, you do this by installing a server and establishing the server as a domain controller in a new forest.

When you set up DNS, be sure you configure the appropriate reverse lookup zones. You should have one reverse lookup zone for each subnet. If you forget to set up the reverse zones and do this after installing your servers, be sure that the appropriate PTR records have been created for your domain controllers and Exchange servers. In Active Directory Sites And Services, check that the sites and subnets are configured appropriately. You need to create a subnet in Active Directory to represent each of the subnets on your network. If DNS reverse zones and Active Directory subnets are not configured properly, you will likely experience long startup times on your servers, and Exchange services will likely not start properly.

Installing Exchange Server

Before you run Exchange Server 2016 Setup, make sure that the server meets the system requirements and prerequisites as discussed in “Navigating Exchange 2016 Editions” in Chapter 13, “Implementing Exchange Services.” You can only run Exchange Server 2016 on full installations of Windows Server 2012, Windows Server 2012 R2 and Windows Server 2016. You cannot install Exchange Server 2016 on a server running in Windows Server Core mode. Instead, you must convert the Core mode to a full installation.

NOTE You can use Setup to install the Exchange Server 2016 management tools on domain-joined computers running 64-bit editions of Windows 8.1 or Windows 10.

You can run Exchange Server 2016 Setup in one of several modes, including:

- **Install** Used when you’re installing a new server role or adding a server role to an existing installation.
- **Upgrade** Used when you have an existing installation of Exchange and you’re installing a service pack or cumulative update.
- **Uninstall** Used when you’re removing the Exchange installation.

IMPORTANT Exchange Server 2016 doesn’t support in-place upgrades from any previous version of Exchange. Further, after you install Exchange Server 2016, you won’t be able to rename the server.

Generally, you should install Exchange Server 2016 on member servers rather than on domain controllers. This will ensure Exchange operates with strictest security allowed and has optimal performance. If you do install Exchange Server 2016 on a domain controller, you won’t be able to demote the server. Once Exchange 2016 is installed, changing a server’s role from a member server to a directory server, or vice versa, isn’t supported.

If something goes wrong with the installation and re-running Setup and following the prompts doesn’t help you resolve the problem, you have several options. You can restore the server from backup or you can run Exchange Server 2016 Setup in recovery mode by running `setup /m:RecoverServer` at a command prompt. If you are recovering to a different server, the server must use the same fully qualified domain name (FQDN) as the failed server.

When you recover a server, you don’t specify the roles to restore. Setup detects the Exchange Server object in Active Directory and installs the corresponding files and configuration automatically. After you recover the server, you can restore databases and reconfigure any additional settings.

When you are ready to run Setup, you can begin the installation and install an Exchange server by completing the following steps:

1. Log on to the server using an administrator account. When you install the Mailbox role, you must use a domain account that is a member of the Enterprise Administrators group. If you've already prepared Active Directory, this account must also be a member of the Exchange Organization Administrators group.

IMPORTANT Before beginning setup, you should close any open Windows PowerShell or Microsoft Management Console (MMC) windows. Otherwise you will see a warning during the readiness checks that you need to close these windows. The installation process makes updates to Windows PowerShell and MMC and typically requires exclusive access.

REAL WORLD Ensure the server's TCP/IP settings are properly configured before beginning setup. Also, ensure that the server is a member of the domain in which you want the Exchange organization to be configured. During setup, the server will try to identify the Active Directory site in which it is located. The server will try to connect with a domain controller and global catalog server in this site.

2. Do one of the following:

- If you are using installation media, insert the Exchange Server 2016 media. If Autorun is enabled, Exchange Server 2016 Setup should start automatically. Otherwise, double-click **Setup.exe** on the root folder of the installation media.
 - If you are using a download, access the folder where you extracted the Exchange setup files and then start Exchange 2016 Setup by double-clicking **Setup.exe**.
3. On the Check For Updates page, you can specify whether to check for updates to the setup process. If you don't want to check for updates, select **Don't Check For Updates Right Now** before you click **Next** to continue. Setup will then copy files and initialize resources. The server also tries to validate the state of Active Directory.

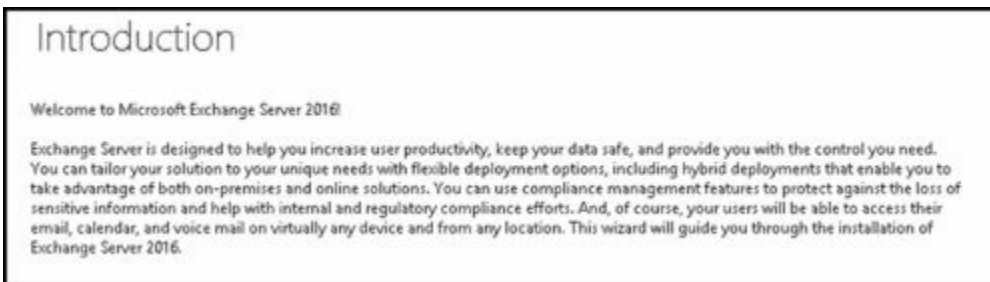


NOTE If the server is unable to validate the state of Active Directory and choose a domain controller to work with, Setup will log errors and may also report that a domain controller could not be located. If errors are reported, do not continue with the installation. Instead, exit Setup and resolve the communication problem.

4. Setup copies files that are required to the server and then prepares resources.



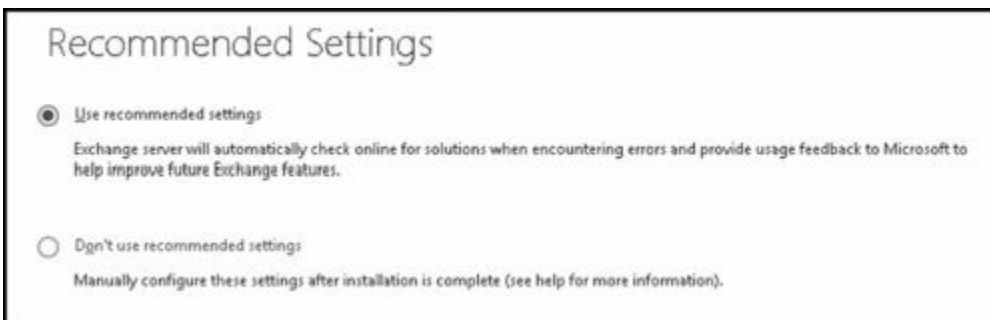
5. The Introduction page begins the installation process. Click **Next** to continue.



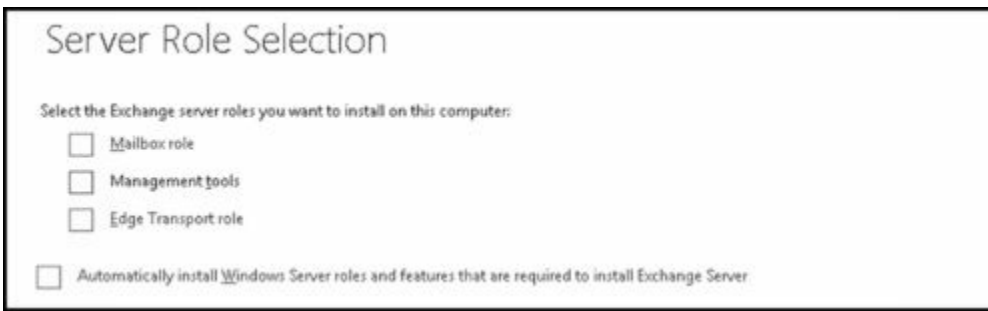
6. On the License Agreement page, review the software license terms. If you agree to the terms, select **I Accept The Terms In The License Agreement**, and then click **Next**.



7. On the Recommended Settings page, select whether you want to use the recommended settings. If you select Use Recommended Settings, Exchange will automatically send error reports and information about your computer hardware and how you use Exchange to Microsoft. If you select Don't Use Recommended Settings, error and usage reporting are disabled but you can enable them at any time after Setup completes. Click **Next** to continue.



8. On the Server Role Selection page, choose whether you want to install the Mailbox role or the Edge Transport role. The management tools are installed automatically if you install any server role.



9. Select **Automatically Install Windows Server Roles And Features That Are Required To Install Exchange Server** to have Setup install any required Windows prerequisites. You may need to reboot the computer to complete the installation of some Windows features. If you don't select this option, you must install the required Windows features manually. Click **Next** to continue.
10. On the Installation Space And Location page, note the space required for the installation. Click **Browse** to choose a location for the installation. Ensure you have enough disk space available on the related drive. Click **Next** to continue.

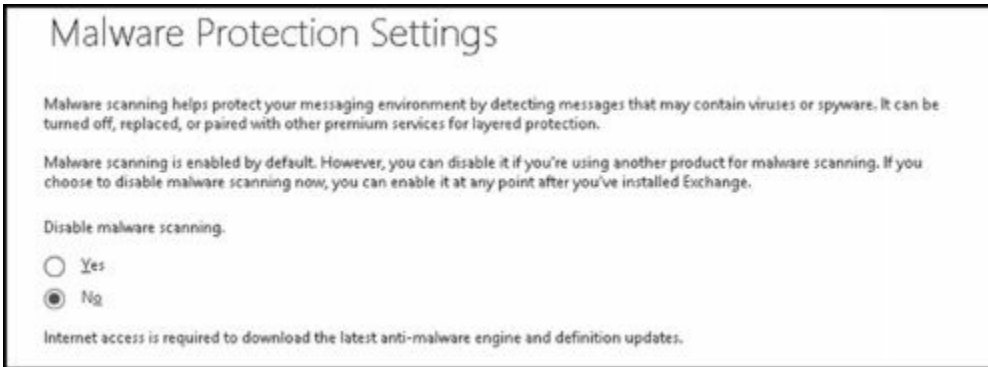


11. If this is the first Exchange server in your organization, on the Exchange Organization page, type a name for your Exchange organization or accept the default value of First Organization. The Exchange organization name must be 64 characters or less and can contain only the characters A through Z, a through z, 0 through 9, space (as long as the space is not leading or trailing), and hyphen or dash. You can't leave the organization name blank. Click **Next** to continue.

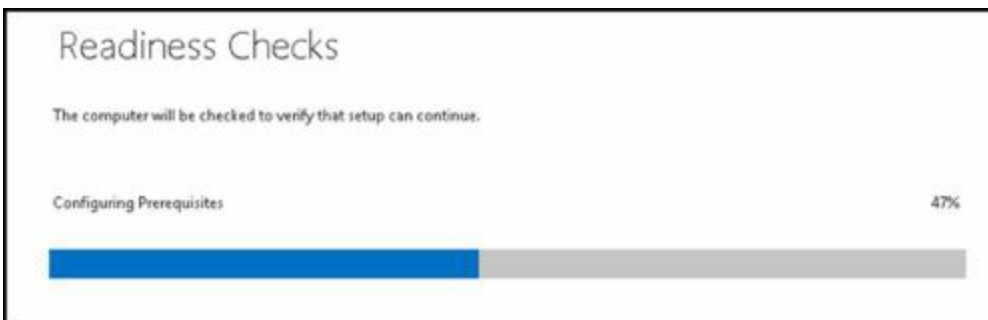
REAL WORLD Exchange 2016 supports shared permissions and split permissions. Split permissions allow organizations to separate Exchange management and Active Directory management. Role Based Access Control (RBAC) are the recommended split permissions model used with Exchange. If you want to use shared permissions or split permissions that use RBAC, do not select the Apply Active Directory Split Permissions... check box. If your organization has strict requirements for separate management of Active Directory and Exchange Server and RBAC will not meet your needs, select the Apply Active Directory Split Permissions... check box. However, you will then be unable to create users, groups, contacts, and other Active Directory objects using the Exchange management tools



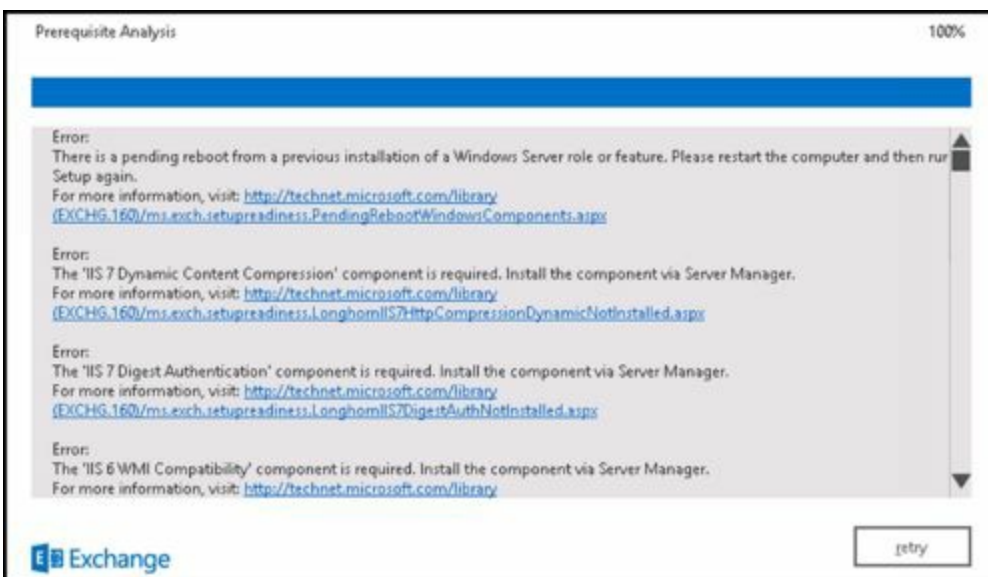
12. If you're installing the Mailbox role, on the Malware Protection Settings page, choose whether you want to enable or disable malware scanning. If you disable malware scanning, it can be enabled later.



13. When you click **Next**, Setup will verify that all prerequisites have been installed and that Exchange 2016 can be installed.

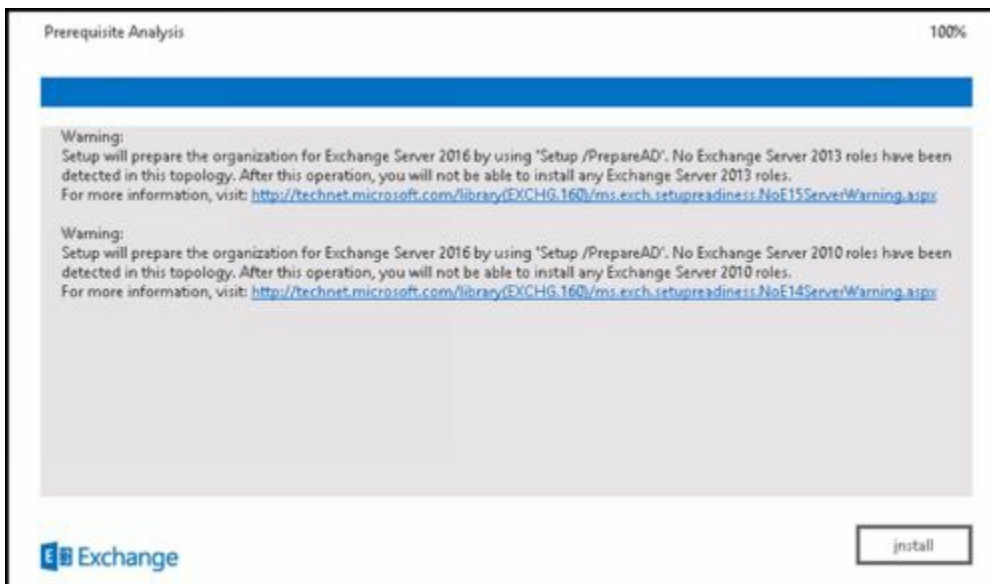


14. Note any errors. You must resolve any reported errors before you can install Exchange Server 2016. For most errors, you don't need to exit Setup. After resolving a reported error, click **Retry** to run the prerequisite checks again.




15. When all readiness checks have completed successfully, note any warnings and then click **Install** to install Exchange 2016. The installation process can take up

to 1 hour.



The Setup Progress page, tracks the progress of the installation. The installation is performed in a series of steps, with the progress for the current step tracked with a progress bar and as a percentage of completion.

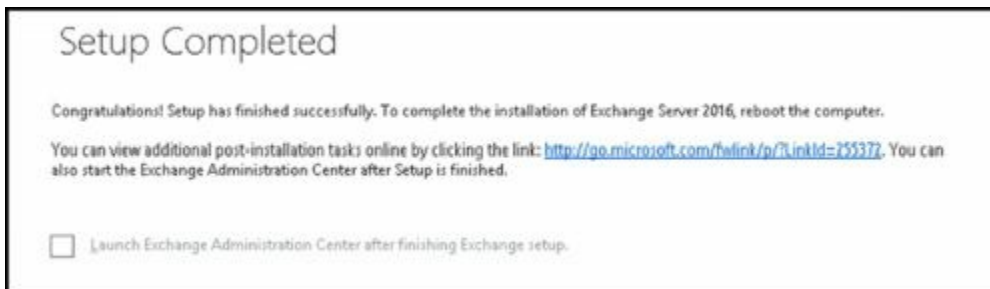


The number of steps varies, depending on the tasks Setup must perform to prepare the environment as well as the options you selected. Typically, the steps see will include:

1. Organization preparation
2. Preparing setup
3. Stopping services
4. Exchange files... Language files
5. Restoring services
6. Languages
7. Management tools
8. Transport service.
9. Client Access service
10. Unified Messaging service
11. Mailbox service
12. Front End Transport service
13. Front End service
14. Finalizing setup

Finally, you'll see the Setup Completed page, when Setup completes the installation. Although you must restart the server to finalize the installation, you may want to select the Launch Exchange Administration Center checkbox before selecting Finish and then set the product key.

NOTE Alternatively, you can manually start the Exchange Administration Center by opening Internet Explorer and entering the Exchange Administration Center. By default, this URL is `https://ServerName/ecp/` where *ServerName* is the name of the server, such as: `https://mailserver35/ecp/`.



By default, Exchange 2016 runs in trial mode. To get out of trial mode, you must validate the installation. In the left pane of the Exchange Admin Center, click Servers and then select Servers as the feature you want to work with. As shown in Figure 15-1, each installed Exchange server is listed by name, role and version. Click the server you want to work with. In the details pane, a link is provided for entering a product key.

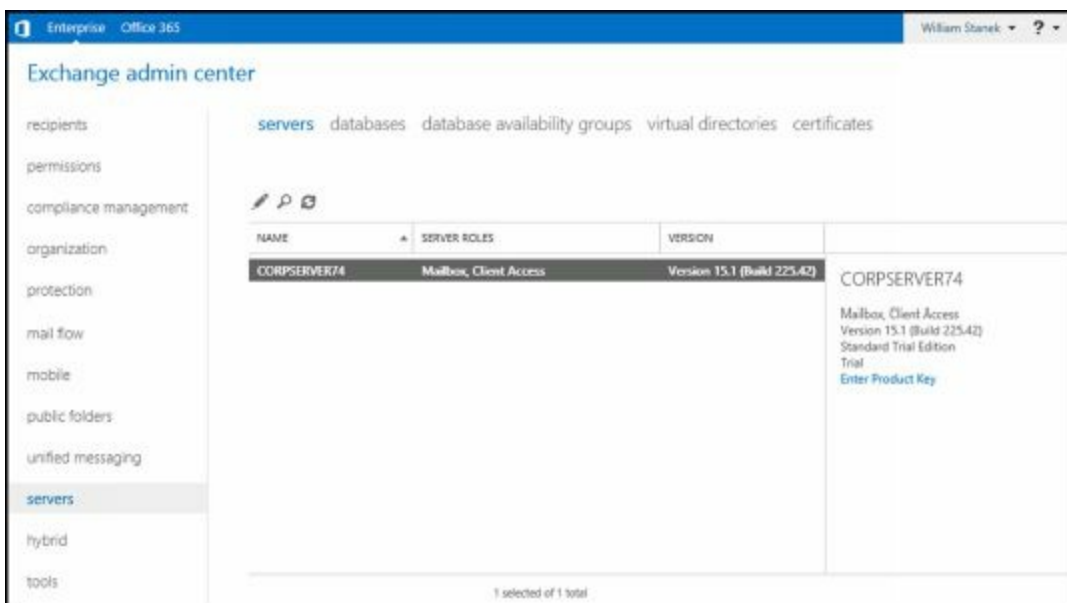


FIGURE 15-1 Viewing installed servers in Exchange Admin Center.

Clicking the Enter Product Key link opens the properties dialog box for the mail server with the General page displayed. Enter a valid product key in the boxes provided and then click Save.

CORPSERVER74

- general
- databases and database availability groups
- POP3
- IMAP4
- unified messaging
- DNS lookups
- transport limits
- transport logs
- Outlook Anywhere

Version number:
Version 15.1 (Build 225.42)

Roles:
Mailbox, Client Access

This Exchange server is currently licensed as a Trial Edition. You can add a new license by entering a product key below. [Learn more](#)

Enter a valid product key.
[] - [] - [] - [] - []

Save Cancel

You can change the product key at any time on the general page. Select Change Product Key, enter a valid product key, and then click Save.

You can upgrade a Standard edition to an Enterprise edition using the options on the General page as well. Select Change Product Key, enter a valid product key for Enterprise edition and then click Save.

Verifying and Completing the Installation

You can verify that Exchange Server 2016 installed successfully by running the `Get-ExchangeServer` cmdlet in Exchange Management Shell. This command displays a list of all Exchange 2016 server roles that are installed on a specified server.

During installation, Exchange Setup logs events in the Application log of Event Viewer. You can review the Application log to make sure there are no warning or error messages related to Exchange setup. Typically, these events have event IDs 1003 and 1004, with the source as `MSExchangeSetup`.

You also can learn more about the installation by reviewing the setup log file created during the setup process. This log file is stored in the `%SystemDrive%\ExchangeSetupLogs` folder with the name `ExchangeSetup.log`. The `%SystemDrive%` variable represents the root directory of the drive where the operating system is installed. Because these logs contain standard text, you can perform a search using the keyword *error* to find any setup errors that occurred.

As discussed previously, Setup must be able to contact Active Directory. If Setup is unable do this, errors will be logged and the Exchange organization will not be prepared properly. In the following example, Setup couldn't validate the state of Active Directory and couldn't locate a domain controller:

```
[02/14/2016 11:05:31.0253] [0] Setup is choosing the domain controller to use
[02/14/2016 11:05:42.0630] [0] Setup is choosing a local domain controller...
[02/14/2016 11:05:45.0033] [0] [ERROR] Setup encountered a problem while
validating the state of Active Directory: Could not find any Domain
Controller in domain imagedlands.com.
[02/14/2016 11:05:45.0158] [0] [ERROR] Could not find any Domain Controller
in domain imagedlands.com.
[02/14/2016 11:05:45.0205] [0] [ERROR] Domain controller not found in the
domain "imagedlands.com".
[02/14/2016 11:05:45.0205] [0] Setup will use the domain controller ".
[02/14/2016 11:05:45.0205] [0] Setup will use the global catalog ".
[02/14/2016 11:05:45.0955] [0] No Exchange configuration container was found for the
organization. Message: 'Could not find any Domain Controller
in domain imagedlands.com.'
```

Because of this problem, Setup didn't fully prepare the organization and had problems configuring the Mailbox role: Transport service and the other services as well. When Setup is able to validate the state of Active Directory, the log records a very different set of events as shown in the following example:

```
[02/14/2016 12:11:07.0115] [0] Setup is choosing the domain controller to use
[02/14/2016 12:11:14.0135] [0] Setup is choosing a local domain controller...
[02/14/2016 12:11:24.0729] [0] Setup has chosen the local domain controller
CorpServer24.imagedlands.com for initial queries
[02/14/2016 12:11:24.0885] [0] PrepareAD has either not been run or has not
```


replicated to the domain controller used by Setup. Setup will attempt to use the Schema Master domain controller CorpServer24.imagedlands.com [02/14/2016 12:11:24.0885] [0] The schema master domain controller is available

[02/14/2016 12:11:24.0901] [0] The schema master domain controller is in the local domain; setup will use CorpServer24.imagedlands.com

[02/14/2016 12:11:24.0901] [0] Setup is choosing a global catalog...

[02/14/2016 12:11:24.0917] [0] Setup has chosen the global catalog server CorpServer24.imagedlands.com.

[02/14/2016 12:11:24.0932] [0] Setup will use the domain controller 'CorpServer24.imagedlands.com'.

[02/14/2016 12:11:24.0932] [0] Setup will use the global catalog 'CorpServer24.imagedlands.com'.

[02/14/2016 12:11:24.0948] [0] No Exchange configuration container was found for the organization. Message: 'Could not find the Exchange Configuration Container'.

Here, Setup was able to select a domain controller to work with, locate the schema master, and choose a global catalog server. Note that Setup reports that PrepareAD was not run or replicated and that no Exchange configuration container was found. This is normal for a new installation of Exchange 2016. Shortly after validating the state of Active Directory, Setup will determine the organization-level operations to perform. For a new installation of Exchange 2016, related entries should look similar to the following:

[02/14/2016 12:11:26.0339] [0] Setup is determining what organization-level operations to perform.

[02/14/2016 12:11:26.0339] [0] Setup has detected a missing value. Setup is adding the value PrepareSchema.

[02/14/2016 12:11:26.0339] [0] Setup has detected a missing value. Setup is adding the value PrepareOrganization.

[02/14/2016 12:11:26.0339] [0] Setup has detected a missing value. Setup is adding the value PrepareDomain.

Here, Setup reports that it will prepare the Active Directory schema, the Exchange organization, and the domain. You can confirm each by looking for the elements that should have been created or configured as discussed in the section titled “Coexistence and Active Directory” earlier in Chapter 14.

To complete the installation for an initial deployment of Exchange into an organization, you need to perform the following tasks:

For Client Access services:

- If you plan to use ActiveSync for mobile messaging clients, configure direct push, authentication, and mobile devices.
- Configure internal and external URLs for the Outlook web applications, Exchange ActiveSync, Exchange Admin Center, and Offline Address Book.
- Configure internal and external URLs for autodiscover, SMTP and HTTP.

- Configure authentication and display options, as appropriate.
- Enable the server for POP3 and IMAP4, as appropriate.
- Enable the server for MAPI over HTTP, as appropriate.
- A self-signed digital certificate is created by default but won't be automatically trusted by clients. You can either establish trust or obtain a certificate from a third party that the client trusts.

For Mailbox servers:

- Configure domains for which you will accept email. You need an accepted domain entry for each SMTP domain for which you will accept email.
- Configure Send connectors as appropriate. If you are unsure about the Send connectors that are needed, create an Internet Send connector at minimum. Use the address space of "*" to route all outbound mail to the Internet.
- If you also deployed the Edge Transport server role, you need to subscribe to the Edge Transport server so that the EdgeSync service can establish one-way replication of recipient and configuration information from Active Directory to the AD LDS store on the Edge Transport server.
- Configure DNS MX resource records for each accepted domain.
- Configure OAB distribution for Outlook 2010 and later clients.
- Configure database availability groups and mailbox database copies, as appropriate.

For Unified Messaging service:

- Configure a unified messaging dial plan, and add the server to it.
- Configure Unified Messaging hunt groups.
- Enable users for unified messaging, as appropriate.
- Configure your IP/VoIP gateways or IP-PBXs to work with Exchange Server.
- Configure a Unified Messaging IP gateway in Exchange Server.
- As desired, create auto-attendant and mailbox policies and configure additional dial plans, gateways, and hunt groups.

Uninstalling Exchange 2016

The Exchange Server 2016 installation process uses Windows Installer. Using Windows Installer helps to streamline and stabilize the installation process, and it makes modification of installation components fairly easy. Thanks to Windows Installer, you also can resume a failed installation or modification simply by re-running Exchange Setup.

If you no longer need an Exchange server, you can uninstall Exchange 2016 to remove Exchange services. Before you can uninstall Exchange 2016, you must disable, move or remove all mailboxes, archive mailboxes, public folder mailboxes and arbitration mailboxes from the server. To help you with this process, use the following techniques:

- Get a list of all mailboxes in all available databases hosted on the server by running the command **Get-MailboxDatabase -Server *ServerName* | Get-Mailbox**. These are the user mailboxes that must be moved or removed.
- Get a list of all archive mailboxes in all available databases hosted on the server by running the command **Get-MailboxDatabase -Server *ServerName* | Get-Mailbox -Archive**. These are the archive mailboxes that must be moved or removed.
- Get a list of all arbitration mailboxes in all available databases hosted on the server by running the command **Get-MailboxDatabase -Server *ServerName* | Get-Mailbox -Arbitration**. These are the arbitration mailboxes that must be moved or removed.
- Disable a non-arbitration mailbox so that you can delete the mailbox database by running the command **Disable-Mailbox *MailboxId***.
- Disable an archive mailbox so you can delete the mailbox database by running the command **Disable-Mailbox *MailboxId* -Archive**.
- Disable a public folder mailbox so that you can delete the mailbox database by running the command **Disable-Mailbox *MailboxId* -PublicFolder**.
- Rather than removing arbitration mailboxes, you should move them to another server using **New-MoveRequest**. If this is the last server in the organization, disable the arbitration mailbox instead by running the command **Disable-Mailbox *ArMailboxID* -Arbitration -DisableLastArbitrationMailboxAllowed**.

Although Exchange Setup doesn't allow you to uninstall a server, you can use Programs And Features to do this. Follow these steps:

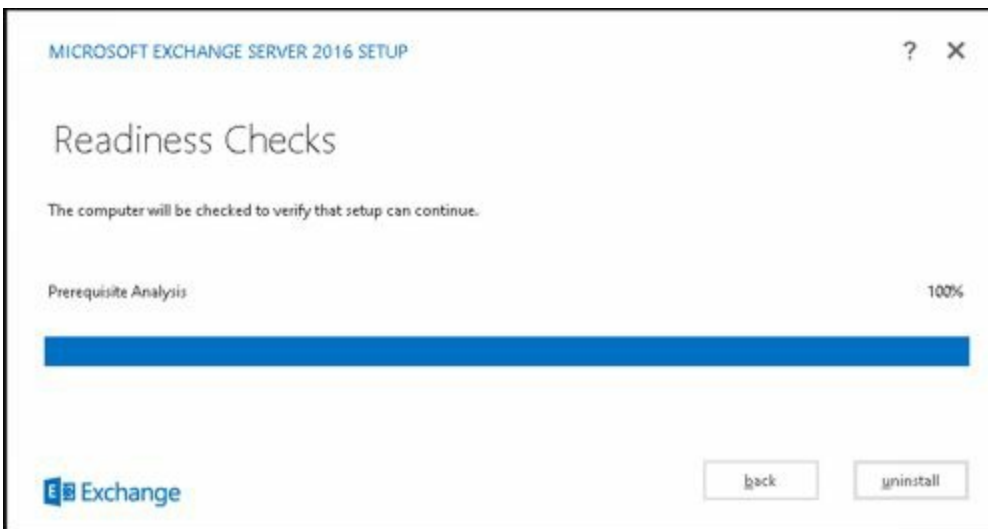
1. Log on to the Exchange server either locally or remotely. Close any local instances of Exchange Management Shell that are open.
2. In Control Panel, click the **Uninstall A Program link** under Programs. In Programs And Features, select the **Microsoft Exchange Server 2016** entry to display the Uninstall button.
3. Click the **Uninstall** button to start Exchange Setup. When Setup finishes initializing, click **Next** to continue.



4. Setup will then perform readiness checks to determine whether Exchange 2016 can be removed from the server. If there are errors, you will need to correct the error conditions. Perform the recommended tasks and then click Retry to have Setup repeat the readiness checks.



5. When Setup completes the readiness checks without error, you can complete the removal process by clicking **Uninstall**.



Using Cumulative Updates

With Exchange Server 2013 and later, Microsoft decided to start delivering routine product updates and security updates separately. Under this servicing model, routine product updates are delivered periodically as a single, cumulative update, and security updates are delivered separately. While this allows you to install security updates as they are released without having to install a cumulative update, cumulative updates themselves will contain security updates. As with earlier releases of service packs in Exchange Server, cumulative updates are delivered as full product updates and installed as an upgrade.

To better align on-premises Exchange and Exchange Online, Microsoft tries to release cumulative updates on a fixed schedule and applies cumulative updates to their hosted Exchange servers prior to official release of an update. Thus, when an update is released you know it has been applied to all Exchange Online servers and all of the mailboxes stored in the cloud.

IMPORTANT Microsoft is releasing cumulative updates for other products, including Lync and SharePoint, on separate fixed schedules as well. Ideally, this will be a quarterly release schedule with four cumulative updates released each year during the product's lifecycle.

What's in Cumulative Updates?

Cumulative updates more closely resemble service packs than rollup updates. Not only may cumulative updates contain hotfixes and security updates, they may also contain new features, product enhancements, and other changes that affect the way the product works. While language modifications were previously limited to Service Pack releases, cumulative updates may contain updates to language resources. A cumulative update also may contain Active Directory schema updates. If so, the schema changes will be additive and backwards compatible with previous release and product versions.

IMPORTANT Cumulative updates do not replace service packs. Microsoft will continue to release service packs for Exchange Server.

Every cumulative update and service pack is a full release of the product. This means, you install cumulative updates and service packs as a product upgrade and that each update package will be larger than the previous product or update package. Because you install cumulative updates and service packs as an upgrade, any customizations you've made to Exchange Server (using web.config files on Mailbox servers, EdgeTransport.exe.config files on Edge Transport servers, registry changes, or other custom configuration options on servers) are not preserved. This means you will lose any customizations. To prevent this, you must save your customizations and then re-apply them after applying a cumulative update or service pack.

REAL WORLD Don't forget that it is possible the upgrade process will fail. If

this happens, you can recover from the failed upgrade like you would recover from a failed service pack installation, which may include running Exchange Server 2016 Setup with a special recovery option. To do this, you enter the command `SETUP /m:RecoverServer`.

In the unlikely event that the upgrade fails and is unrecoverable, you will need to re-install Exchange Server. This re-installation process will create a new server object and should not result in the loss of mailbox or queue data. However, you will need to re-seed or re-attach existing databases after the re-installation process.

How Are Cumulative Updates Applied?

You apply cumulative updates and service packs using Exchange Server Setup. Because each cumulative update and service pack is a new build of Exchange Server 2016, you don't need to apply cumulative updates or service packs in sequence. You can apply the latest cumulative update or service pack at any time. For example, if you deployed Exchange Server 2016 RTM but didn't upgrade to Cumulative Update 1, you could upgrade the original installation directly to Cumulative Update 2.

IMPORTANT When you are deploying Exchange servers, you don't need to deploy Exchange Server 2016 RTM and then upgrade to a cumulative update or service pack later. As each cumulative update or service pack is a complete build, you can fully deploy the Exchange server using only the current cumulative update or service pack.

In a Database Availability Group configuration, all servers should be running the same cumulative update or service pack of Exchange Server 2016—except during an upgrade. During an upgrade, individual servers within a Database Availability Group can have different cumulative update or service pack versions. This mixed state is expected to be only temporary. Database Availability Group should not operate in a mixed state for long periods of time.

Cumulative updates and service packs are published at the Microsoft Download Center. Because staying current with cumulative updates and service packs may present a particular challenge for some Exchange installations, it is important to note that cumulative updates are supported only for three months after the release of the subsequent cumulative update. With Microsoft's goal of delivering cumulative updates quarterly, this typically means that a prior cumulative update is supported for about six months.

How Do I Track Exchange Version Numbers?

Versioning with Exchange Server 2016 gets a little tricky. The official release of Exchange Server 2016 is referred to as Exchange Server 2016 RTM. Cumulative updates for this release are referred to using the full release name plus the cumulative update number. Thus, Exchange Server 2016 with Cumulative Update 1 is referred to as

Exchange Server 2016 CU1.

As Microsoft releases service packs for Exchange Server 2016, those service packs will be full product rollups that include prior cumulative updates of the product and become part of the release cycle in place of a particular cumulative update. This is why Exchange 2013 had a CU1, CU2, CU3 and then an SP1, followed by a CU5.

Keep in mind the version of Exchange Server is updated when you install a cumulative update or service pack. This means that one way to determine what cumulative update or service pack is applied is to check the version number of the Exchange server. The build number for Exchange 2016 RTM is 225.42 and the build number is incremented for each cumulative update and service pack. You can determine the Exchange version number by entering either of the following commands:

```
Get-Command ExSetup.exe | % {$_.FileVersionInfo}
```

```
Get-ExchangeServer | Format-List Name, Edition, AdminDisplayVersion
```

Installing Cumulative Updates and Service Packs

As discussed previously, cumulative updates and service packs are full builds of Exchange Server 2016. You install a cumulative update or service pack as an upgrade, and there is no rollback process should installation fail. Because of this, you should ensure you have a full recovery plan in place prior to applying a cumulative update. Typically, this means having server backups and other backup plans in place prior to installing an update.

You'll find cumulative updates and service packs for Exchange Server 2016 on the Microsoft Download Center. A single download is provided for both Exchange Server 2016 Enterprise and Exchange Server 2016 Standard. When you access the download page, click Download to start the download process. Next, copy the download to your computer for installation at a later time by clicking Save. Copy the download to your server if necessary.

When you run the executable, Windows verifies the file, and you'll then be able to extract the download to a folder. Be sure to specify an exact folder to put all the setup components in one place. Within this folder, you'll find a program called Setup.exe. This is the Exchange Server 2016 Setup program.

Preparing to Install a Cumulative Update or Service Pack

Before you run Exchange Setup make sure you read the release notes for the cumulative update or service pack. Also make sure that any server on which you plan to install the cumulative update or service pack meets the system requirements and prerequisites for Exchange Server 2016.

You can run Exchange Server 2016 only on full installations of Windows Server 2012 and later. Exchange Server 2016 doesn't support in-place upgrades from any previous version of Exchange. After you install a cumulative update or service pack, you cannot uninstall the cumulative update or service pack to revert to an earlier version of Exchange Server 2016. If you uninstall a cumulative update or service pack, Exchange Server 2016 is removed from the server.

As cumulative updates and service packs may contain Active Directory schema changes and other Active Directory updates, you may want to update Active Directory prior to deploying a cumulative update or service pack on any server in your organization, especially in a large enterprise. Here, keep the following in mind:

- If the update contains schema changes, run the following command prior to executing the Exchange Server 2016 Setup.exe: **setup.exe /PrepareSchema /IAcceptExchangeServerLicenseTerms.**
- If the update contains enterprise Active Directory changes (such as role-based Access Control updates), run the following command prior to executing the Exchange Server 2016 Setup.exe: **setup.exe /PrepareAD /IAcceptExchangeServerLicenseTerms.**

- If the update contains changes to the permissions within the Active Directory domain partition, run **setup.exe /PrepareDomain /IAcceptExchangeServerLicenseTerms** in each domain containing Exchange servers or mailboxes.

Although Exchange Server 2016 Setup can perform these processes for you during the upgrade, the changes can take some time to replicate throughout a large organization. By performing these tasks manually, you can streamline the upgrade process. You also can ensure the tasks are run with accounts that have appropriate permissions. Keep the following in mind:

- If schema needs to be updated and you haven't previously prepared schema, you must ensure the account you use is delegated membership in the Schema Admins group.
- If you're installing the first Exchange 2016 server in the organization, the account you use must have membership in the Enterprise Admins group.
- If you've already prepared the schema and aren't installing the first Exchange 2016 server in the organization, the account you use must be a member of the Exchange 2016 Organization Management role group.

NOTE Administrators who are members of the Delegated Setup group can deploy Exchange 2016 servers that have been previously provisioned by a member of the Organization Management group.

After you install a cumulative update or service pack, you must restart the server so that changes can be made to the registry and operating system. If something goes wrong with the installation and re-running Setup and following the prompts doesn't help you resolve the problem, you have several options. You can restore the server from backup or you can run Exchange Server 2016 Setup in recovery mode by running **setup /m:RecoverServer** at a command prompt. If you are recovering to a different server, the server must use the same FQDN as the failed server.

When you recover a server, Setup detects the Exchange Server object in Active Directory and installs the corresponding files and configuration automatically. After you recover the server, you can restore databases and reconfigure any additional settings.

Installing a Cumulative Update or Service Pack

When you are ready to run Setup and install an update, you can begin the installation. If you are installing a new server using a current cumulative update or service pack, follow the procedure as discussed previously under "Installing Exchange Server." Otherwise, to update an existing installation of Exchange 2016, complete the following steps:

1. Log on to the server using an administrator account. When you install the Mailbox role, you must use a domain account that is a member of the Enterprise Administrators group and a member of the Exchange Organization Administrators group.

IMPORTANT Before beginning setup, you should close any open Windows

PowerShell or MMC windows. Otherwise you will see a warning during the readiness checks that you need to close these windows. The installation process may make updates to Windows PowerShell and MMC which requires exclusive access.

CAUTION If you are applying a cumulative update or service pack to an existing Exchange 2016 server, any customized per-server settings you made in Exchange configuration files will be overwritten. To prevent this, save your customized settings before you run Setup. This will help you easily re-configure your server after the update.

2. Access the folder where you extracted the Exchange setup files and then start Exchange 2016 Setup by double-clicking **Setup.exe** . If you've enabled User Access Control (UAC), you must right-click Setup.exe and select Run As Administrator.
3. On the Check For Updates page, you can specify whether to check for updates to the setup process. If you don't want to check for updates, select **Don't Check For Updates** before you click **Next** to continue. Setup will then copy files and initialize resources.

The server also tries to validate the state of Active Directory. If the server is unable to validate the state of Active Directory and choose a domain controller to work with, Setup will log errors and may also report that a domain controller could not be located. If errors are reported, do not continue with the installation. Instead, exit Setup and resolve the communication problem.

4. If you are installing a new server, you'll see the Introduction page. If you are updating an existing server, you'll see the Upgrade page. Click **Next** to continue.

IMPORTANT Seeing the Upgrade page is a confirmation that Setup identified the existing Exchange 2016 installation on the server. There is a problem if you are applying an update or service pack to a server already running Exchange 2016 and don't see the Upgrade page at this point. You may need to restart the server or resume Exchange services that have been stopped and then re-run Setup.

5. On the License Agreement page, review the software license terms. If you agree to the terms, select **I Accept The Terms In The License Agreement** , and then click **Next** .
6. On the Readiness Checks page, ensure the prerequisite checks completed successfully. If they haven't, you must resolve any reported errors before you can update Exchange Server 2016. For most errors, you don't need to exit Setup. After resolving a reported error, click **Retry** to run the prerequisite checks again.

NOTE A cumulative update or service pack may require additional Windows components.

7. When all readiness checks have completed successfully, click **Install** to update

Exchange 2016. The installation process should take about 60 minutes.

The Setup Progress page tracks the progress of the installation. The installation is performed in a series of steps, with the progress for the current step tracked with a progress bar and as a percentage of completion. The number of steps varies, depending on the tasks Setup must perform to prepare the environment as well as the options you selected.

As part of the update, Setup removes existing Exchange files from the installation and then copies new files into the appropriate directories. Finally, you'll see the Setup Completed page, when Setup completes the installation.

You must restart the server to finalize the installation. You can verify that the update to Exchange 2016 installed successfully by running the following command to confirm the version has been updated:

```
Get-ExchangeServer | Format-List Name, Edition, AdminDisplayVersion
```

During installation, Exchange Setup logs events in the Application log. You can review the Application log to make sure there are no warning or error messages related to Exchange setup. You also can learn more about the installation by reviewing the setup log file created during the setup process. This log file is stored in the %SystemDrive%\ExchangeSetupLogs folder with the name ExchangeSetup.log. The %SystemDrive% variable represents the root directory of the drive where the operating system is installed. Review any error entries in the log file.

Chapter 16. Exchange 2016 Administration Essentials

To work effectively with Exchange 2016, you need to master a number of key concepts, including:

- [How to access and work with Exchange Admin Center](#)
- [How connections are authenticated and proxied](#)
- [How Exchange uses virtual directories](#)
- [Why Exchange requires SSL certificates](#)
- [Which Windows processes are used with Exchange Server](#)

For troubleshooting or deeper work with Exchange, you also need to know how to bypass Exchange Admin Center and Exchange Management Shell so that you can work directly with Exchange Server. These topics are all covered in this chapter.

Working with Exchange Admin Center

Whether you are working with on-premises, online, and hybrid Exchange organizations, you'll use one graphical tool for administration: Exchange Admin Center. Exchange Admin Center is a browser-based application that you access via the Mailbox servers deployed in your Exchange organization. This application can be configured with an internal access URL and an external access URL. However, as only an internal access URL is configured by default, you can initially only access Exchange Admin Center when you are on the corporate network. To access this application when working outside the organization, you'll need to configure an external access URL.

Accessing Exchange Admin Center

Although Exchange Admin Center replaces Exchange Management Console and Exchange Control Panel (ECP), ECP continues to be the name for the related virtual directory. You access Exchange Admin Center by following these steps:

1. Exchange Admin Center works with most browsers, as long as you are using a current version. In your browser, enter the secure URL for Exchange Admin Center. If you are outside the corporate network, enter the external URL, such as <https://mail.tvpress.com/ecp> . If you are inside the corporate network, enter the internal URL, such as <https://mailserver48/ecp>.

The version of Exchange Admin Center you see depends on the version of Exchange running on the Mailbox server hosting your personal mailbox. Exchange 2016 runs version 15, and you can specify this version explicitly by appending ? **ExchClientVer=15** to the internal or external URL. For example, if your external URL is <https://mail.tvpress.com>, you could enter **<https://mail.tvpress.com/ecp?ExchClientVer=15>** as the URL.

NOTE By default, you must use HTTPS to connect. If you don't, you'll see an error stating "Access is denied." Using HTTPS ensures data transmitted between the client browser and the server is encrypted and secured.

2. If your browser displays a security alert stating there's a problem with the site's security certificate or that the connection is untrusted, proceed anyway. This alert is displayed because the browser does not trust the self-signed certificate that was automatically created when the Exchange server was installed.
- With Internet Explorer, the error states "There's a problem with this website's security certificate." Proceed by selecting the Continue To This Web Site (Not Recommended) link.
 - With Google Chrome, the error states "Your connection is not private." Continue by clicking the Advanced link and then clicking the Proceed To ... link.
 - With Mozilla Firefox, the error states "This connection is untrusted." Proceed by selecting I Understand The Risks and then selecting Add Exception. Finally, in the

Add Security Exception dialog box, select Confirm Security Exception.

3. You'll see the logon page for Exchange Admin Center. Enter your user name and password, and then click **Sign In**.

Be sure to specify your user name in DOMAIN\username format. The domain can either be the DNS domain, such as imaginedlands.com, or the NetBIOS domain name, such as imaginedlands. For example, the user JoeH could specify his logon name as imaginedlands.com\joeH or imaginedlands\joeH.

4. If you are logging on for the first time, select your preferred display language and time zone, and then click **Save**.

After you log in to Exchange Admin Center, you'll see the list view with manageable features listed in the Navigation menu (see Figure 16-1). When you select a feature in the Navigation menu, you'll then see the related topics or "tabs" for that feature. The manageable items for a selected topic or tab are displayed in the main area of the browser window. For example, when you select Recipients in the Navigation menu, the topics or tabs that you can work with include: Mailboxes, Groups, Resources, Contacts, Shared and Migration.

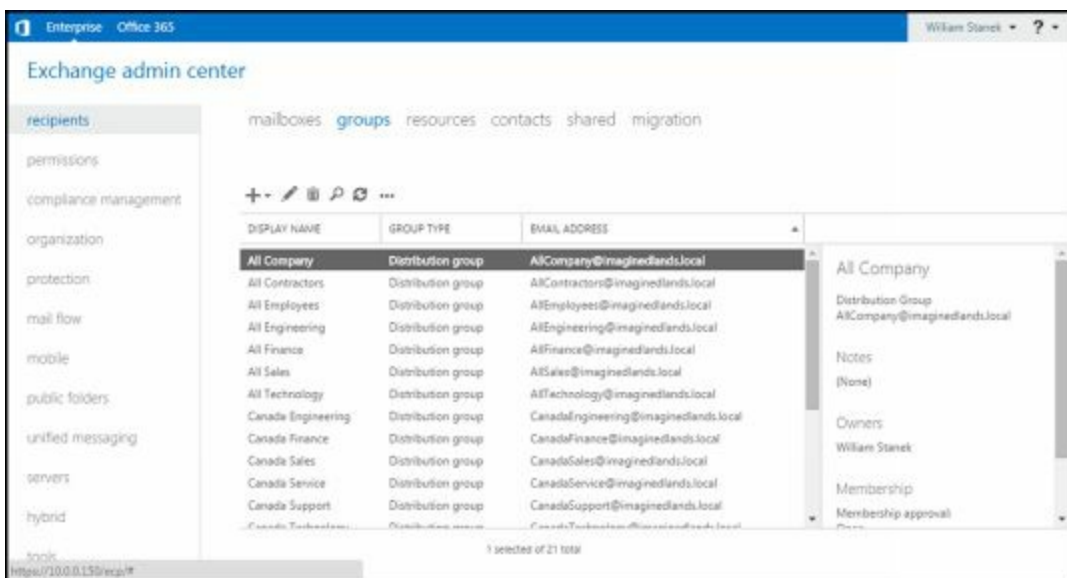


FIGURE 16-1 Exchange Admin Center uses a list view with manageable features listed on the left.

Working with Exchange Server Certificates

When you install an Exchange server, the setup process creates several self-signed security certificates that are used for authentication. The default certificates available for Mailbox servers include:

- **Microsoft Exchange** A self-signed certificate used by IMAP, POP, IIS, and SMTP. If Autodiscover is configured, this certificate is also used for Autodiscover. This is the primary certificate used by Exchange.
- **Microsoft Server Auth Certificate** A self-signed certificate for authenticating SMTP connections.
- **Exchange Delegation Federation** A self-signed certificate used when federated

sharing is configured in the Exchange organization.

- **WMSVC** A self-signed certificate used by the Windows Management service.

As Figure 16-2 shows, you can view these certificates in Exchange Admin Center by selecting Servers in the Navigation menu and then selecting Certificates. Because the default certificates are not issued by a trusted authority, you see a related error message whenever you use HTTPS to access services hosted by your Mailbox servers, including Exchange Admin Center, the PowerShell application, and Microsoft Outlook Web App.

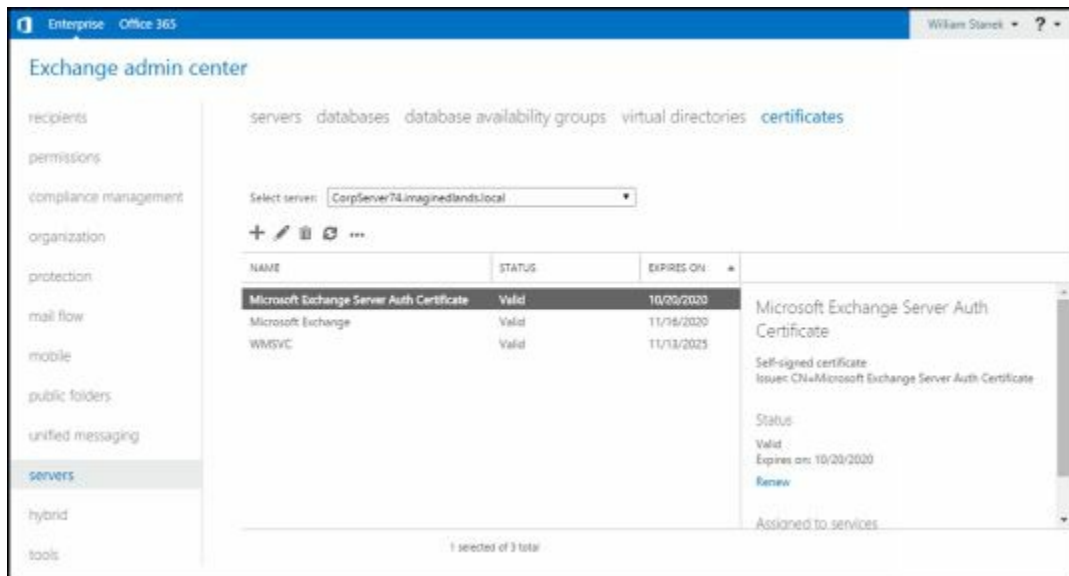


FIGURE 16-2 Viewing the SSL certificates installed on Exchange servers.

One way to eliminate this error message is to install a certificate from a trusted authority on your Mailbox servers. Web browsers should already be configured to trust certificates issued by your organization's certification authority (CA) or by a trusted third-party authority. Typically, browsers need additional configuration only when you use your own CA with non-domain-joined machines.

The services a certificate can be used with include Internet Message Access Protocol (IMAP), Post Office Protocol (POP), SMTP, Internet Information Services (IIS), and Unified Messaging (UM). The default self-signed certificates are assigned services automatically during setup based on the roles installed on the Exchange server.

When you work with certificates, it's critical that you ensure the certificate is used for the right subject name and alternative names. As an example, the Microsoft Exchange certificate created by default has the Subject set as `cn=ServerName`, where `ServerName` is the name of the server, such as `cn=MailServer21`, and the Subject Alternative Names is set as `DNS Name=ServerName`, `DNS NAME=FullyQualifiedServerName`, and `DNS Name=DomainName`. If Autodiscover is configured, there's also a Subject Alternative Name entry for `DNS Name=Autodiscover.DomainName`. For example, `MailServer21` in the `Imagedlands.com` domain means the subject name is set as:

`cn=MailServer21`

and the Subject Alternative Name entries typically are:

`DNS Name = MailServer21`

DNS Name = MailServer21.imagedlands.com

DNS Name = imagedlands.com

DNS Name = Autodiscover.imagedlands.com

REAL WORLD Outlook Web App (OWA) and Exchange Admin Center can become inaccessible as a result of a required SSL certificate becoming corrupted or being invalidated. If this happens, you will need to access Exchange directly and re-create the required certificate or certificates. One way to safeguard yourself against this problem is to create copies of the original certificates using the Certificates snap-in. When you add this snap-in to a Microsoft Management Console, specify that you want to manage certificates for a computer account. You'll then find the certificates under the Personal node. Export each certificate in turn using the Certificate Export Wizard. To start this wizard, right-click a certificate, select All Tasks, and then select Export.

If your organization has a CA, have your security administrator issue a certificate. Generate the certificate by completing the following steps.

1. In a web browser, open Certificate Services by entering the appropriate URL, such as <https://CertServer03/certsrv>.
2. Specify that you want to create a new request and then choose the advanced creation option.
3. Submit a certificate request by using a base 64 encoded PKS #7 or PKS #12 file.
4. Once the certificate request file is generated, open the file in a text editor.
5. While you are working with Certificate Services in your browser, access the request. Copy the contents of the certificate request file and paste them into the request.
6. Select web server as the server type, and leave all other attributes blank.
7. Save the certificate.

After you create the certificate, you must make it available on the designated Exchange server. To do this, access the Exchange server and then import the certificate using `Import-ExchangeCertificate`. Next, use `Enable-ExchangeCertificate` to enable the certificate for specific Exchange services.

If you can purchase a certificate from a trusted third-party authority, you also must make the certificate available on the designated Exchange server. To do this, access the Exchange server and then import the certificate using `Import-ExchangeCertificate`. Next, use `Enable-ExchangeCertificate` to enable the certificate for specific Exchange services. Finally, ensure that the new certificate is in use and test web services using `Test-OutlookWebServices` as shown in the following example:

```
test-outlookwebservices | fl
```

By default `Test-OutlookWebServices` verifies the Availability service, Outlook Anywhere, Offline Address Book, and Exchange Web Services. Test MAPI over HTTP using `-ProbeIdentity OutlookMapiHttpSelfTestProbe` with `Test-OutlookConnectivity`.

Test connectivity to the OWA and ECP virtual directories using Test-OwaConnectivity and Test-EcpConnectivity respectively. However, before you can use any of the Test cmdlets, you must create a test account by running the Scripts\New-TestCasConnectivityUser.ps1 script. You'll find this script in the %ExchangeInstallPath%, which by default is C:\Program Files\Microsoft\Exchange Server\V15\. The password you set for the test account is temporary and will be automatically changed every 7 days.

Once you've imported and enabled the certificate, you can then view the certificate in Exchange Admin Center or by using Get-ExchangeCertificate to confirm it is configured as expected. You'll want to ensure the status is valid, the expiration date is appropriate, the subject name is correct, the subject alternative names are correct, and that the assigned services are appropriate.

Configuring Exchange Admin Center

You can configure Exchange Admin Center for single-server and multiserver environments. In a single-server environment, you use one Mailbox server for all of your remote management needs. In a multiple-server environment, you can instruct administrators to use different URLs to access different Mailbox servers, or you can use load-balancing and give all administrators the same access URL.

IMPORTANT You should work out a plan for Exchange namespaces as part of your deployment and configuration. You will need a namespace for autodiscover, HTTP, SMTP and other technologies you've deployed in the Exchange organization, such as IMAP or POP3. If you have a large Exchange organization, each Exchange namespace can be load balanced across your datacenters in a layer 7 configuration that doesn't rely on session affinity. With a datacenter pair, you could achieve an approximately 50/50 traffic split between the datacenters using round robin DNS or a similar technique.

NOTE You can use Exchange Admin Center with firewalls. You configure your network to use a perimeter network with firewalls in front of the designated Mailbox servers and then open port 443 to the IP addresses of your Mailbox servers. If Secure Sockets Layer (SSL) is enabled and you want to use SSL exclusively, you only need port 443, and you don't need to open port 80.

You can manage the Exchange Admin Center application using Internet Information Services (IIS) Manager or Exchange Management Shell. The related commands for Exchange Management Shell are as follows:

- **Get-ECPVirtualDirectory** Displays information about the ECP application running on the Web server providing services for Exchange. By default, only front-end virtual directories are listed. Add `-ShowMailboxVirtualDirectories` to also display the back-end virtual directories.

```
Get-ECPVirtualDirectory [-Identity AppName ]  
[-ADPropertiesOnly <$true | $false>]
```

`[-ShowMailboxVirtualDirectories <$true | $false>]`

`[-DomainController DomainControllerName]`

`Get-ECPVirtualDirectory -Server ExchangeServerName`

`[-ADPropertiesOnly <$true | $false>]`

`[-ShowMailboxVirtualDirectories <$true | $false>]`

`[-DomainController DomainControllerName]`

- **New-ECPVirtualDirectory** Creates a new ECP application running on the Web server providing services for Exchange. You should use this command only for troubleshooting scenarios where you are required to remove and re-create the ECP virtual directory.

`New-ECPVirtualDirectory [-AppPoolId AppPoolName]`

`[-DomainController DomainControllerName] [-ExternalUrl URL]`

`[-InternalUrl URL] [-WebSiteName SiteName]`

- **Remove-ECPVirtualDirectory** Use the `Remove-ECPVirtualDirectory` cmdlet to remove a specified ECP application providing services for Exchange.

`Remove-ECPVirtualDirectory -Identity AppName`

`[-DomainController DomainControllerName]`

- **Set-ECPVirtualDirectory** Modifies the configuration settings for a specified ECP application providing services for Exchange. Set `-AdminEnabled` to `$false` to turn off Internet access to the Exchange Admin Center.

`Set-ECPVirtualDirectory -Identity AppName`

`[-AdminEnabled <$true | $false>]`

`[-BasicAuthentication <$true | $false>] [-DomainController`

`DomainControllerName] [-ExternalAuthenticationMethods Methods]`

`[-DigestAuthentication <$true | $false>]`

`[-FormsAuthentication <$true | $false>]`

`[-ExternalUrl URL] [-GzipLevel <Off | Low | High | Error>]`

`[-InternalUrl URL] [-LiveIdAuthentication <$true | $false>]`

`[-WindowsAuthentication <$true | $false>]`

- **Test-ECPConnectivity** Displays information about the ECP application running on the Web server providing services for Exchange.

`Test-ECPConnectivity [-ClientAccessServer ServerName]`

`[-MailboxServer ServerName] [-DomainController DomainControllerName]`

`[-RTSEndPoint EndPointID] [-TestType <Internal | External>]`

`[-MonitoringContext <$true | $false>]`

`[-ResetTestAccountCredentials <$true | $false>]`

`[-Timeout NumSeconds] [-TrustAnySSLCertificate <$true | $false>]`

`[-VirtualDirectoryName DirectoryName]`

At Exchange Management Shell prompt, you can confirm the location of the Exchange Admin Center application by typing `get-ecpvirtualdirectory` .

`Get-ECPVirtualDirectory` lists the name of the application, the associated website, and the server on which the application is running, as shown in the following example:

Name	Server
-----	-----
ecp (Default Web Site)	MailServer62

In this example, a standard configuration is being used in which the application named ECP is running on the Default Web Site on MailServer62. You can use `Set-ECPVirtualDirectory` to specify the internal and external URL to use as well as the permitted authentication types. Authentication types you can enable or disable include basic authentication, Windows authentication, and Live ID basic authentication. You can use `New-ECPVirtualDirectory` to create or re-create an ECP application on a Web server providing services for Exchange and `Remove-ECPVirtualDirectory` to remove an ECP application. You can verify that Exchange Admin Center is working properly using `Test-ECPCConnectivity`.

The PowerShell application has a similar set of commands. In Exchange Management Shell, the related commands are `New-PowerShellVirtualDirectory`, `Get-PowerShellVirtualDirectory`, `Set-PowerShellVirtualDirectory`, and `Test-PowerShellConnectivity`. If you enter **`Get-PowerShellVirtualDirectory | Format-List`**, you'll get configuration details for each Exchange server in the Exchange organization. You can use `Set-PowerShellVirtualDirectory` to enable or disable authentication mechanisms, including basic authentication, certificate authentication, Live ID basic authentication, Live ID NTLM negotiate authentication, and Windows authentication. You can also specify the internal and external URLs for the PowerShell virtual directory on a per-server basis. By default, servers have only internal URLs for PowerShell. For troubleshooting issues related to the PowerShell virtual directory, enter **`Test-PowerShellConnectivity`** followed by the URL to test, such as `https://mailer1.tvpress.com/powershell`.

You'll also find commands for working with virtual directories related to:

- Outlook Web App, including `New-OwaVirtualDirectory`, `Get-OwaVirtualDirectory`, `Set-OwaVirtualDirectory`, and `Remove-OwaVirtualDirectory`
- Offline Address Books, including `New-OabVirtualDirectory`, `Get-OabVirtualDirectory`, `Set-OabVirtualDirectory`, and `Remove-OabVirtualDirectory`
- Autodiscover, including `New-AutodiscoverVirtualDirectory`, `Get-AutodiscoverVirtualDirectory`, `Set-AutodiscoverVirtualDirectory`, and `Remove-AutodiscoverVirtualDirectory`
- Exchange ActiveSync, including `New-ActiveSyncVirtualDirectory`, `Get-ActiveSyncVirtualDirectory`, `Set-ActiveSyncVirtualDirectory`, and `Remove-ActiveSyncVirtualDirectory`
- MAPI over HTTP, including `New-MapiVirtualDirectory`, `Get-MapiVirtualDirectory`, `Set-MapiVirtualDirectory`, and `Remove-MapiVirtualDirectory`

Keep in mind that there are separate but interconnected virtual directories representing front-end and back-end services. Typically, front-end virtual directories are used for authentication and proxying while back-end virtual directories are used for actual processing. Although the front-end and back-end virtual directories have different

components and configurations, the Exchange cmdlets for creating these virtual directories are designed to configure the appropriate settings and components for either front-end or back-end use as appropriate.

You should specify explicitly whether you want to work with the front-end or back-end components. You do this by specifying the related website name. The Default Web Site is used by the front-end components and the Exchange Back End website is used by back-end components.

Bypassing Exchange Admin Center and Troubleshooting

Exchange makes extensive use of IIS. Front-end apps for each essential client service are configured on the Default Web Site. Back-end apps for each essential client service are configured on the Exchange Back End website. This means you'll find front-end and back-end apps for ActiveSync, Autodiscover, ECP, EWS, MAPI, OAB, PowerShell and RPC connections made over HTTP.

Understanding Remote Execution in Exchange Admin Center

When you access OWA in a web browser, you are performing remote operations via the PowerShell application running on the Web server providing Exchange services whether you are logged on locally to an Exchange server or working remotely. The same is true for ECP, but the process is a little more complex, as shown in the following high-level view of the login and workflow process:

1. Generally, OWA handles the initial login for ECP. Thus, when you access ECP using a URL such as <https://mailserver17/ecp>, the browser actually is redirected to OWA with a URL such as <https://mailserver17/owa/auth/logon.aspx?replaceCurrent=1&url=https%3a%2f%2fmailserver17%2fecp%2f>.
2. Once you log in to Exchange, you are connected to the designated Mailbox server using the ECP app running on the Default Web Site.
3. ECP performs authentication checks that validate your access to the Exchange 2016 server and determine the Exchange role groups and roles your account is a member of. You must be a member of at least one management role.
4. ECP creates a remote session with the Exchange 2016 server. A remote session is a runspace that establishes a common working environment for executing commands on remote computers.
5. By default, you are connected to the Mailbox server on which your user mailbox resides with the front-end ECP app on that server acting as proxy for the backend ECP app on that server.
6. As you perform tasks, these tasks are executed via the PowerShell app, which also has front-end and back-end components.

IMPORTANT Every step of the login and workflow process relies on properly configured SSL certificates. HTTPS uses SSL certificates to establish and encrypt connections. SSL certificates are also used to initialize and validate remote sessions. Although you could disable the requirement for HTTPS and allow HTTP to be used for connections, the remote sessions themselves would still rely on properly configured SSL certificates.

Thus, many interconnected components must be functioning correctly for you to connect to and work with Exchange Server.

Bypassing Exchange Admin Center and Exchange Management Shell

Because Exchange Management Shell uses remote sessions that run via the PowerShell application running on IIS, you often need a way to work directly with Exchange Server, especially when you are trying to diagnose and resolve problems. Intuitively, you might think that you should do this in the same way you establish a remote session with Exchange Online. For example, if you want to connect to Server21, you might want to use the following code:

```
$Session = New-PSSession -ConfigurationName Microsoft.Exchange  
-ConnectionUri https://server21/powershell/ -Authentication Basic  
-Credential wrstanek@imaginedlands.com -AllowRedirection
```

```
Import-PSSession $Session
```

However, if there are any configuration problems, including issues with SSL certificates, you won't be able to connect to or work with Exchange Server in this way. Instead, you'll have to bypass the web-based management interfaces and connect directly to Exchange server using the following technique:

1. Log on to the Mailbox server you want to work with—either at the console or using a remote desktop connection.
2. Open an administrative PowerShell window by right-clicking **Windows PowerShell** and then clicking **Run As Administrator**.
3. Import all Exchange-related snapins for Windows PowerShell by entering **Add-PSSnapin *exchange***. You'll then be able to work directly with Exchange and any related cmdlets.

Because Exchange has separate front-end and back-end components, you'll often need to perform troubleshooting tasks on both. Rather than log on locally to a server, you may want to work remotely. You can invoke commands, establish direct remote sessions, or execute commands remotely using the `-ComputerName` parameter available with certain cmdlets. (For more information, see Chapters 4 and 5 in *Windows PowerShell: The Personal Trainer*. [Stanek & Associates, 2015])

To invoke commands on remote servers or establish a direct remote session, use the following technique:

1. Log on to any workstation or server where you've installed the Exchange management tools. (Doing so ensures the Exchange related snap-ins are available.)
2. Open an administrative PowerShell window by right-clicking **Windows PowerShell**, and then clicking **Run As Administrator**.

3. Import all Exchange-related snapins for Windows PowerShell by entering **Add-PSSnapin *exchange*** .
4. Either invoke commands on the remote Exchange server or establish a remote session with the remote Exchange server. In your remote sessions, be sure to connect directly, as shown in the following example:

```
$Session = New-PSSession -computername mailsrv26  
-Credential imaginedlands\williams @techjob
```

```
Import-PSSession $Session
```

IMPORTANT When you work with Exchange in this way, you establish connections via the Windows Remote Management (WinRM) service. On an Exchange server, WinRM and related services are set up automatically. On your management computer, you may need to install the required components and configure WinRM as discussed previously in “Using Exchange Management Shell” in Chapter 13, “Implementing Exchange Services.” See also “Customizing Remote Management Services” later in this chapter.

Troubleshooting OWA, ECP, Powershell, and More

Sometimes users and administrators see a blank page or an error when they try to log in to OWA or ECP. This problem and other connection issues, such as those related to OAB, Autodiscover, and PowerShell, can occur because of a wide variety of configuration issues, including:

- [Invalid or missing TCP/IP settings](#)
- [Corrupted or improperly configured virtual directories](#)
- [Missing, expired, invalid, or improperly configured SSL certificates](#)

However, before you look at specific issues, ensure required services are running as discussed in “Maintaining Exchange Services” later in this chapter. Be sure to examine the services related to both front-end and back-end services.

Typically, the next logical step is to validate the TCP/IP settings of your Mailbox servers. Not only do servers need to communicate with each other and clients, they also need to communicate with domain controllers.

If Exchange Server can’t communicate properly with a domain controller, you may see an error similar to the following when you open Exchange Admin Center or Exchange Management Shell:

The LDAP server is unavailable.

Description: An unhandled exception occurred during the execution of the current web request. Please review the stack trace for more information about the error and where it originated in the code.

Exception Details: System.DirectoryServices.Protocols.LdapException: The LDAP server is unavailable.

Source Error:

An unhandled exception was generated during the execution of the current web request. Information regarding the origin and location of the exception can be identified using the exception stack trace below.

Stack Trace:

```
[LdapException: The LDAP server is unavailable.]
  System.DirectoryServices.Protocols.LdapConnection.Connect() +160015
  System.DirectoryServices.Protocols.LdapConnection.BindHelper
(NetworkCredential newCredential, Boolean needSetCredential) +264
  Microsoft.Exchange.Data.Directory.PooledLdapConnection.BindWithRetry
(Int32 maxRetries) +702
```

Resolve the problem by doing the following:

- Ensure the server has the proper TCP/IP settings and is connected to the network.
- Ensure a domain controller is available for the server to communicate with.

Users or administrators may see a blank page when they try to log on to OWA or ECP as a result of a configuration or certificate problem. If you've determined that required services are running and that the TCP/IP settings are correct, next try to isolate and identify the specific issue.

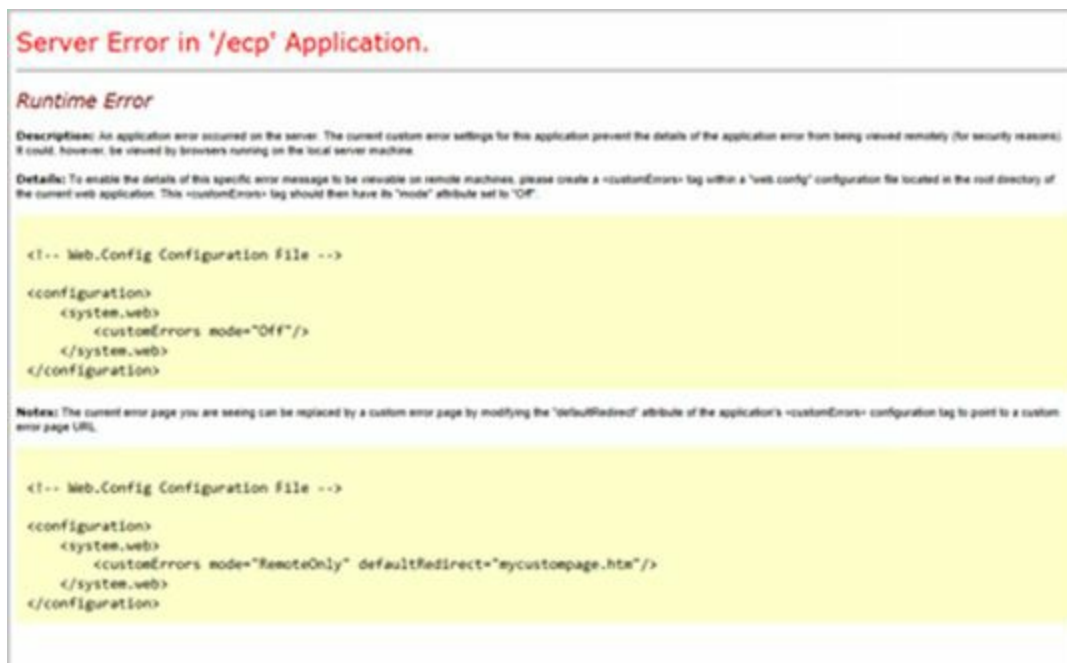


FIGURE 16-3 An error related to an improperly configured virtual directory or a misconfiguration in IIS.

Try to log on to OWA or ECP in a browser. Sometimes when you log on to OWA or ECP, you'll see a runtime error that indicates an improperly configured virtual directory or an application error due to misconfiguration in IIS (see Figure 16-3). Other times, the browser window may simply be empty or blank as mentioned previously.

For deeper troubleshooting, log on to the Mailbox server hosting the mailbox for the users or administrators experiencing the problem and open Exchange Management Shell. If there's a problem with SSL certificates rather than virtual directory configuration, you'll see an error similar to the following:

```
New-PSSession : [mailserver08] Connecting to remote server mailserver08
failed with the following error message : The server certificate on the
destination computer (mailserver08:443) has the following errors:
The SSL certificate is signed by an unknown certificate authority. For more
information, see the about_Remote_Troubleshooting Help topic.
```

```
At line:1 char:12
```

```
+ $Session = New-PSSession -ConfigurationName Microsoft.Exchange
-ConnectionUri http ...
```

```
+ ~~~~~
~~~~~
```

```
+ CategoryInfo          : OpenError
(System.Manageme....RemoteRunspace:RemoteRunspace) [New-PSSession],
PSRemotingTransportException
+ FullyQualifiedErrorId : 12175,PSSessionOpenFailed
```

If there's a problem with virtual directory configuration, you may see another type of error, such as:

```
New-PSSession : [mailserver55.imagedlands.com] Processing data from
remote server mailserver55.imagedlands.com failed with the following
error message: The WinRM Shell client cannot process the request. The shell
handle passed to the WSMAN Shell function is not valid. The shell handle is
valid only when WSMANCreateShell function completes successfully. Change
the request including a valid shell handle and try again. For more
information, see the about_Remote_Troubleshooting Help topic.
```

```
At line:1 char:1
```

```
+ New-PSSession -ConnectionURI "$connectionUri" -ConfigurationName
Microsoft.Excha ... + ~~~~~
```

```
~~~~~
```

```
+ CategoryInfo          : OpenError:
(System.Manageme....RemoteRunspace:RemoteRunspace) [New-PSSession],
PSRemotingTransportException
+ FullyQualifiedErrorId : -2144108212,PSSessionOpenFailed
```

To help diagnose the problem, you can test services using Test-OutlookWebServices. By default, Test-OutlookWebServices verifies the Availability service, Outlook Anywhere, Offline Address Book, and Unified Messaging. You can test OWA, ECP, and PowerShell using Test-OwaConnectivity, Test-EcpConnectivity, and Test-PowerShellConnectivity respectively.

Resolving SSL Certificate Issues

To resolve a certificate issue, you'll need to restore or recreate the primary SSL certificate on the Mailbox server. By default, the self-signed certificate named Microsoft

Exchange is the certificate used for authentication and encrypting communications whenever you use OWA, ECP, or the management tools to work with Exchange. If you backed up the certificates on the server or exported the certificates as discussed previously in this chapter in “Working with Exchange Server certificates,” you can restore the original certificate to restore services.

If you don't have a backup or an export of the primary SSL certificate, you'll need to re-create the certificate. You can create a new self-signed certificate using `New-ExchangeCertificate`. The following example shows how to configure services, the subject name, and subject alternative names for MailServer08 in the Imaginedlands.com domain:

```
New-ExchangeCertificate -SubjectName "cn=MailServer08"  
-DomainName imaginedlands.com -IncludeServerFQDN  
-Services IIS, IMAP, POP, SMTP
```

IMPORTANT If there's a problem preventing you from using Exchange Admin Center and Exchange Management Shell, you'll need to bypass the web-based management interfaces and connect directly to Exchange Server using the technique discussed earlier in the chapter.

With certificates issued by a local CA or a third-party CA, you can use the original certificate file. Import the certificate using `Import-ExchangeCertificate` and then use `Enable-ExchangeCertificate` to enable the certificate for IIS, IMAP, POP, and SMTP services. You can ensure that the certificate is in use and test services as discussed previously.

Resolving OWA, ECP, or Other Virtual Directory Issues

To resolve a virtual directory issue, you can remove and then re-create the virtual directory. You won't always know whether the problem exists in the front-end configuration, the back-end configuration, or both, so you may need to remove and re-create both virtual directories. I recommend removing and re-creating the front-end virtual directory first and then checking to see if this resolves the problem before removing and re-creating the back-end virtual directory.

As an example, if you've determined the OWA virtual directory is misconfigured, you can remove it using `Remove-OwaVirtualDirectory` and then re-create it using `New-OwaVirtualDirectory`. For example, the following commands remove and then re-create the OWA virtual directory from the Default Web Site on MailServer08:

```
remove-owavirtualdirectory -identity "mailserver08\owa (Default Web Site)"
```

```
new-owavirtualdirectory -server mailserver08  
-websitename "Default Web Site"
```

IMPORTANT Keep in mind that if there's a problem preventing you from using Exchange Admin Center and Exchange Management Shell, you'll need to bypass the web-based management interfaces and connect directly to Exchange Server

using the technique discussed earlier in the chapter. You'll then be able to remove the virtual directory and then re-create it. When you are logged on to the server you are configuring, you don't need to use the `-Server` parameter with `New-OwaVirtualDirectory`.

By default, the `New-OwaVirtualDirectory` and `New-EcpVirtualDirectory` commands enable basic authentication and forms authentication but do not enable Windows authentication. Because Windows authentication is required for OWA and ECP, you must use the `Set-OwaVirtualDirectory` and `Set-EcpVirtualDirectory` commands to modify the default authentication settings. The following example enables Windows authentication and disables basic and forms authentication:

```
set-owavirtualdirectory -identity "mailserver08\owa (Default Web Site)"  
-WindowsAuthentication $True -BasicAuthentication $false  
-FormsAuthentication $false
```

After you re-create a virtual directory you should restart IIS services. You can do this in IIS Manager or by entering the following command at an elevated command prompt or shell:

```
iisreset
```

You can then test the service using `Test-OwaConnectivity`, or you can try to log in to OWA. If this doesn't resolve the problem, you can remove, re-create, and configure the OWA virtual directory for the back-end services, as shown in this example:

```
remove-owavirtualdirectory -identity "mailserver55\owa (Exchange Back End)"
```

```
new-owavirtualdirectory -server mailserver55  
-websitename "Exchange Back End"
```

```
set-owavirtualdirectory -identity "mailserver55\owa (Exchange Back End)"  
-WindowsAuthentication $True -BasicAuthentication $false  
-FormsAuthentication $false
```

Complete the process by restarting IIS services and then check to ensure the problem is resolved. If the problem isn't resolved, look to related services. For example, remote PowerShell must be properly configured for OWA and ECP to work. If you suspect the PowerShell virtual directory is misconfigured, you can remove and re-create it as well.

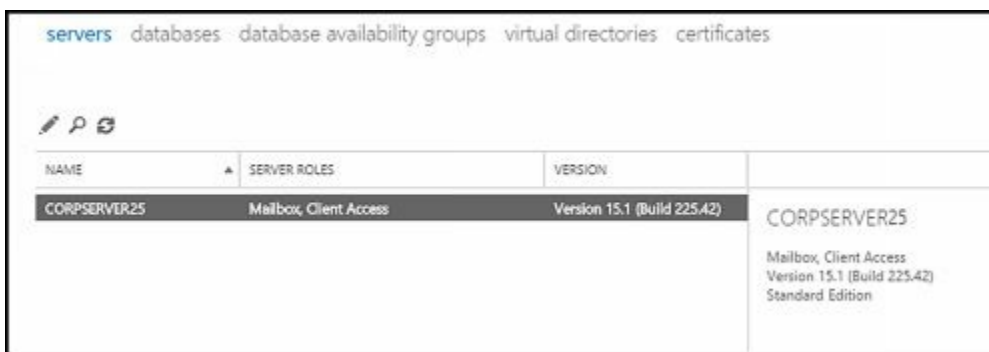
Validating Exchange Server Licensing

With Exchange Server 2016, you do not enter a product key during initial setup. Instead, you provide the product key after installation using Exchange Admin Center. Until you enter a product key, Exchange Server 2016 runs in trial mode.

The product key you provide determines which edition is established on an Exchange server. You can use a valid product key to go from a trial edition to Standard Edition or Enterprise Edition of Exchange Server 2016 without having to reinstall the program.

To determine the established edition and licensing for an Exchange server complete the following steps:

1. In Exchange Admin Center, select **Servers** in the Navigation menu.
2. In the main pane, select the server you want to work with.
3. Look in the Details pane to see the server roles, version, established edition, and license details.

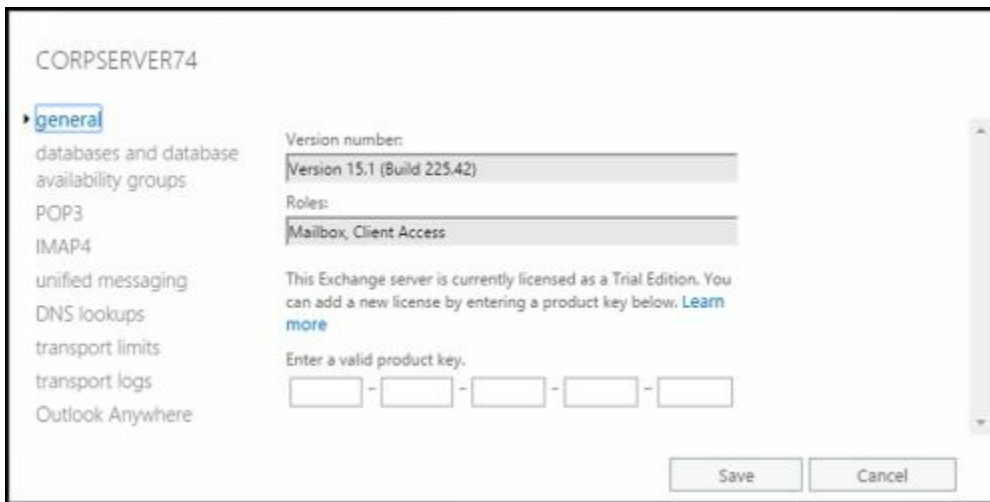


To enter a product key complete the following steps:

1. In Exchange Admin Center, select **Servers** in the Navigation menu.



2. In the main pane, select the server you want to work with.
3. In the details pane, select **Enter Product Key**. This opens the Exchange Server dialog box.



4. Enter the product key for the Exchange Server 2016 edition you want to establish, either Standard or Enterprise, and then click **Save** .

NOTE The product key is a 25-character alphanumeric string, grouped in sets of five characters separated by hyphens. You can find the product key on the Exchange Server 2016 media or license.

5. You should see a dialog box stating the product key has been validated and the product ID has been created. If there's a problem with the product key, you'll see an invalid key warning. Click **OK** . Re-enter or correct the product key and then click Save again. Keep the following in mind:
 - Whenever you set or change the product key on a Mailbox server, you must restart the Microsoft Exchange Information Store service to apply the change. In PowerShell, you can do this by entering **restart-service msexchangeis**.
 - While you can upgrade from Standard to Enterprise edition simply by entering a key for Enterprise edition, you cannot use product keys to downgrade editions. To downgrade editions, you must uninstall Exchange Server and then reinstall the older version.

Using Exchange Management Shell, you can enter a server's product key using the Set-ExchangeServer cmdlet. Sample 16-1 shows the syntax and usage. For the identity parameter, use the server's name, such as MailServer55.

SAMPLE 16-1 Setting the Exchange product key syntax and usage

Syntax

```
Set-ExchangeServer -Identity ' ServerName '
-ProductKey ' ProductKey '
```

Usage

```
Set-ExchangeServer -Identity 'MailServer55'
-ProductKey 'AAAAA-BBBBBB-CCCCC-DDDDD-EEEEEE'
```

TIP By using a valid product key, you can change from the Standard to the Enterprise edition. You also can relicense an Exchange server by entering a new product key for the installed edition, which is useful if you accidentally used the

same product key on multiple servers and want to correct the mistake. The best way to do this is to enter the product key using the Set-ExchangeServer cmdlet.

Using and Managing Exchange Services

Each Exchange server in the organization relies on a set of services for routing messages, processing transactions, replicating data, and much more. “Services for Exchange Server” in Chapter 13, “Implementing Exchange Services” lists these services.

TIP Of all the Exchange services, the one service that relies on having a network connection at startup is the Microsoft Exchange Information Store service. If you start an Exchange server and the server doesn’t have a network connection, the Microsoft Exchange Information Store service might fail to start. As a result, you might have to manually start the service. Sometimes, you’ll find the service has a Stopping state. In this case, you have to wait until the server completely stops the service before you restart it.

Working with Exchange Services

To manage Exchange services, use the Services node in the Computer Management console, which you start by completing the following steps:

1. Start Computer Management. Type **compmgmt** in the Search box, and then select **Computer Management** . Or, on the Tools menu in Server Manager, select Computer Management.
2. To connect to a remote Exchange server, right-click the Computer Management entry in the console tree, and then select **Connect To Another Computer** from the shortcut menu. You can now choose the Exchange server for which you want to manage services.
3. Expand the Services And Applications node, and then select **Services** .

Figure 16-4 shows the Services view in the Computer Management console. The key fields of this window are as follows:

- **Name** The name of the service.
- **Description** A short description of the service and its purpose.
- **Status** The status of the service as started, paused, or stopped. (Stopped is indicated by a blank entry.)
- **Startup Type** The startup setting for the service.

NOTE Automatic services are started when the computer is started. Manual services are started by users or other services. Disabled services are turned off and can’t be started. To start a disabled service, you must first enable it and then start it.

- **Log On As** The account the service logs on as. The default, in most cases, is the local system account.

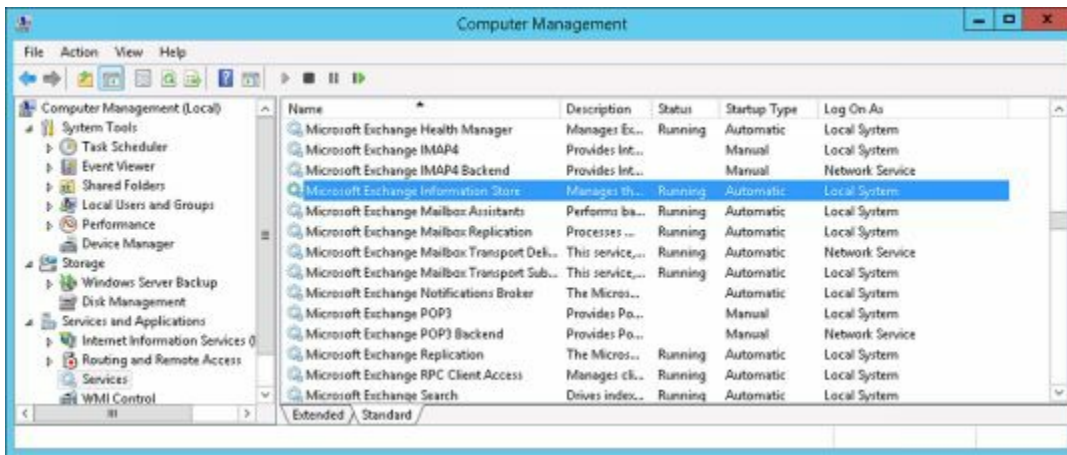


FIGURE 16-4 Using the Services node of the Computer Management console to manage Exchange Server services.

Checking Required Services

You can use Test-ServiceHealth to determine whether all Windows services that Exchange requires are running. As shown in the following example and sample output, the command output lists required services that are running as well as required services that aren't running for each configured Exchange role:

test-servicehealth

```
Role                : Mailbox Server Role
RequiredServicesRunning : True
ServicesRunning      : {IISAdmin, MExchangeADTopology, MExchangeDelivery,
MExchangeIS, MExchangeMailboxAssistants, MExchangeRepl, MExchangeRPC,
MExchangeServiceHost, MExchangeSubmission, MExchangeThrottling,
MExchangeTransportLogSearch, W3Svc, WinRM}
ServicesNotRunning   : {}
```

```
Role                : Client Access Server Role
RequiredServicesRunning : True
ServicesRunning      : {IISAdmin, MExchangeADTopology, MExchangeIMAP4,
MExchangeMailboxReplication, MExchangePOP3, MExchangeRPC,
MExchangeServiceHost, W3Svc, WinRM}
ServicesNotRunning   : {}
```

```
Role                : Unified Messaging Server Role
RequiredServicesRunning : True
ServicesRunning      : {IISAdmin, MExchangeADTopology,
MExchangeServiceHost, MExchangeUM, W3Svc, WinRM}
ServicesNotRunning   : {}
```

```
Role                : Hub Transport Server Role
RequiredServicesRunning : True
ServicesRunning      : {IISAdmin, MExchangeADTopology, MExchangeEdgeSync,
MExchangeServiceHost, MExchangeTransport, MExchangeTransportLogSearch,
```


W3Svc, WinRM}
ServicesNotRunning : {}

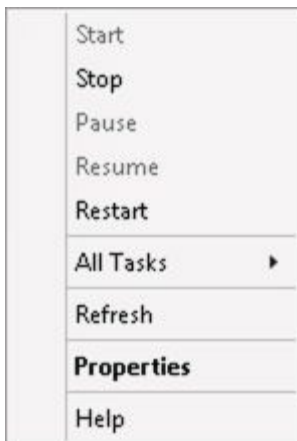
IMPORTANT Although these roles are no longer separate, Mailbox servers perform each role and continue to view each role as separate though interconnected. If there's a problem preventing you from using Exchange Admin Center and Exchange Management Shell, you'll need to bypass the web-based management interfaces and connect directly to Exchange Server using the technique discussed earlier in the chapter.

Maintaining Exchange Services

As an administrator, you'll often have to start, stop, or pause Exchange services. You manage Exchange services through the Computer Management console or through the Services console.

To start, stop, or pause services in the Computer Management console, follow these steps:

1. If necessary, connect to the remote Exchange server for which you want to manage services, as discussed earlier in this section.
2. Expand the Services And Applications node, and then select **Services**.
3. Right-click the service you want to manipulate, and then select Start, Stop, or Pause, as appropriate. You can also choose Restart to have Windows stop and then start the service after a brief pause. Also, if you pause a service, use the Resume option to resume normal operation.



TIP When services that are set to start automatically fail, the status is listed as blank, and you usually receive notification in a pop-up window. Service failures can also be logged to the system's event logs. You can configure recovery actions to handle service failure automatically. For example, you can have Windows attempt to restart the service for you. See the section of this chapter titled "Configuring Service Recovery" for details.

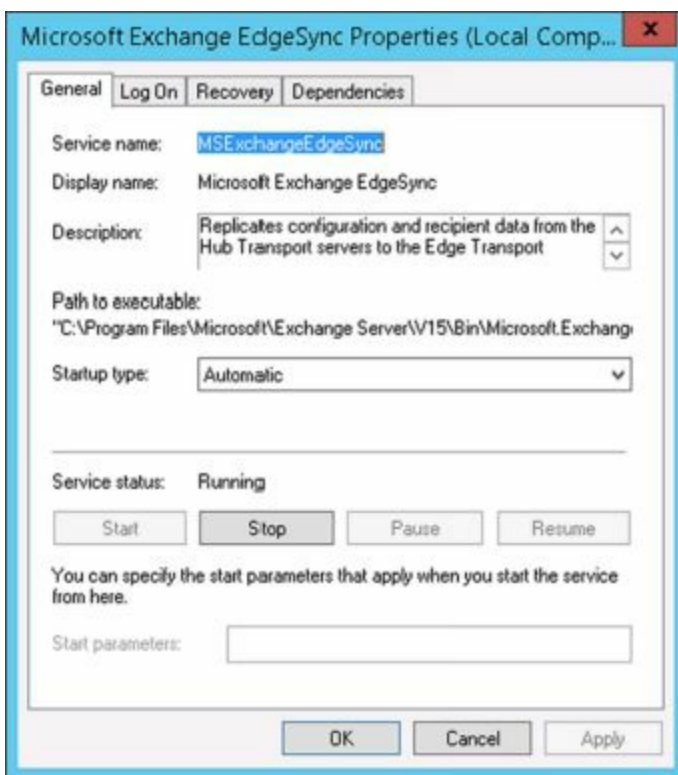
Configuring Service Startup

Essential Exchange services are configured to start automatically and normally shouldn't

be configured with another startup option. That said, if you're troubleshooting a problem, you might want a service to start manually or you might want to temporarily disable a service.

Configure service startup by completing the following steps:

1. In the Computer Management console, connect to the Exchange server for which you want to manage services.
2. Expand the Services And Applications node, and then select **Services** .
3. Right-click the service you want to configure, and then select **Properties** .
4. On the General tab, use the Startup Type drop-down list to choose a startup option. Select Automatic to start a service when the computer starts. Select Manual to allow services to be started manually. Select Disabled to disable the service. Click **OK** .



NOTE The Disabled option doesn't stop the service if it's currently running. It just prevents the service from starting the next time you start the server. To stop the service, you must click Stop.

Configuring Service Recovery

You can configure Windows services to take specific actions when a service fails. For example, you can attempt to restart the service or reboot the server. To configure recovery options for a service, follow these steps:

1. In the Computer Management console, connect to the computer for which you want to manage services.
2. Expand the Services And Applications node, and then select **Services** .

3. Right-click the service you want to configure, and then select **Properties** .



4. On the Recovery tab, you can configure recovery options for the first, second, and subsequent recovery attempts. The available options are as follows:

- Take No Action
- Restart The Service
- Run A Program
- Restart The Computer

5. Configure other options based on your previously selected recovery options. If you elected to restart the service, you need to specify the restart delay. After stopping the service, Windows Server waits for the specified delay period before trying to start the service. In most cases, a delay of one to two minutes should be sufficient. Click **OK**.

When you configure recovery options for critical services, you might try to restart the service on the first and second attempts and then reboot the server on the third attempt. If you notice that a service keeps failing, do some troubleshooting to diagnose and resolve the underlying issue causing the failure.

Customizing Remote Management Services

The Exchange management tools use the Microsoft .NET Framework, Windows Remote Management (WinRM), and Windows PowerShell for remote management. WinRM is implemented in the Windows Remote Management service, which is also referred to as the WS-Management Service or simply the Management Service. To remotely manage Exchange, your management computer must run this service and be configured to use the transports, ports, and authentication methods that your Exchange servers use. The Exchange server you want to connect to must also run this service. If this service isn't

running on your management computer and on the server, remote connections will fail. For remote management, you normally connect to the PowerShell virtual directory configured in IIS on a Mailbox server.

By default, the Management Service connects to and listens on TCP port 80 for HTTP connections and on TCP port 443 for secure HTTP connections. Because firewalls and proxy servers might affect your ability to connect to remote locations over these ports, talk with your company's network or security administrator to determine what steps need to be taken to allow administration over these ports. Typically, the network/security administrator will have to open these TCP ports to allow remote communication between your computer or network and the remote server or network.

The Management Service is preconfigured to share ports with IIS when it runs on the same computer, but it does not depend on IIS. To support remote management, you need to install basic authentication and Windows authentication for IIS on your Exchange servers. These authentication techniques are used when you work remotely.

When you are working with an elevated, administrator command prompt, you can use the WinRM command-line utility to view and manage the remote management configuration. Type **winrm get winrm/config** to display detailed information about the remote management configuration. As Listing 16-1 shows, this lists the configuration details for every aspect of WinRM.

LISTING 16-1 Sample configuration for WinRM

```
Config
MaxEnvelopeSizekb = 500
MaxTimeoutms = 60000
MaxBatchItems = 32000
MaxProviderRequests = 4294967295
Client
  NetworkDelaysms = 5000
  URLPrefix = wsman
  AllowUnencrypted = false
Auth
  Basic = true
  Digest = true
  Kerberos = true
  Negotiate = true
  Certificate = true
  CredSSP = false
DefaultPorts
  HTTP = 5985
  HTTPS = 5986
TrustedHosts
Service
  RootSDDL = O:NSG:BAD:P(A;;GA;;;BA)(A;;GR;;;IU)S:P(AU;
FA;GA;;;WD)(AU;SA;GXGW;;;WD)
  MaxConcurrentOperations = 4294967295
```

```
MaxConcurrentOperationsPerUser = 1500
EnumerationTimeoutms = 240000
MaxConnections = 300
MaxPacketRetrievalTimeSeconds = 120
AllowUnencrypted = false
Auth
  Basic = false
  Kerberos = true
  Negotiate = true
  Certificate = false
  CredSSP = false
  CbtHardeningLevel = Relaxed
DefaultPorts
  HTTP = 5985
  HTTPS = 5986
IPv4Filter = *
IPv6Filter = *
EnableCompatibilityHttpListener = false
EnableCompatibilityHttpsListener = false
CertificateThumbprint
  AllowRemoteAccess = true
```

Winrs

```
AllowRemoteShellAccess = true
IdleTimeout = 7200000
MaxConcurrentUsers = 10
MaxShellRunTime = 2147483647
MaxProcessesPerShell = 25
MaxMemoryPerShellMB = 1024
MaxShellsPerUser = 30
```

If you examine the listing, you'll notice there is a hierarchy of information. The base of this hierarchy, the Config level, is referenced with the path `winrm/config`. Then there are sublevels for client, service, and WinRS, referenced as `winrm/config/client`, `winrm/config/service`, and `winrm/config/winrs`, respectively. You can change the value of most configuration parameters by using the following command:

```
winrm set ConfigPath @{ ParameterName = " Value " }
```

where *ConfigPath* is the configuration path, *ParameterName* is the name of the parameter you want to work with, and *Value* sets the value for the parameter, such as:

```
winrm set winrm/config/winrs @{MaxShellsPerUser="4"}
```

In this example, the `MaxShellsPerUser` parameter is set under `WinRM/Config/WinRS`. Keep in mind that some parameters are read-only and cannot be set in this way.

WinRM requires at least one listener to indicate the transports and IP addresses on which management requests can be accepted. The transport must be HTTP, HTTPS, or both. With HTTP, messages can be encrypted only using NTLM or Kerberos encryption. With HTTPS, Secure Sockets Layer (SSL) is used for encryption. You can examine the

configured listeners by typing **winrm enumerate winrm/config/listener** . As Listing 16-2 shows, this lists the configuration details for configured listeners.

LISTING 16-2 Sample configuration for listeners

Listener

```
Address = *
Transport = HTTP
Port = 5985
Hostname
Enabled = true
URLPrefix = wsman
CertificateThumbprint
ListeningOn = 192.168.1.225, 127.0.0.1, ::1, fe80::5efe:10.0.0.150%13
```

By default, your computer is likely to be configured to listen on any IP address. If so, you won't see any output. To limit WinRM to specific IP addresses, the computer's local loopback address (127.0.0.1) and assigned IPv4 and IPv6 addresses can be explicitly configured for listening. You can configure a computer to listen for requests on HTTP on all configured IP addresses by typing:

```
winrm create winrm/config/listener?Address=*+Transport=HTTP
```

You can listen for requests on HTTPS on all IP addresses configured on the computer by typing:

```
winrm create winrm/config/listener?Address=*+Transport=HTTPS
```

In this case, the * indicates all configured IP addresses. Note that the CertificateThumbprint property must be empty for the SSL configuration to be shared with another service.

You can enable or disable a listener for a specific IP address by typing:

```
winrm set winrm/config/listener?Address=IP:192.168.1.225+Transport=HTTP
@{Enabled="true"}
```

or

```
winrm set winrm/config/listener?Address=IP:192.168.1.225+Transport=HTTP
@{Enabled="false"}
```

You can enable or disable basic authentication on the client by typing:

```
winrm set winrm/config/client/auth @{Basic="true"}
```

or

```
winrm set winrm/config/client/auth @{Basic="false"}
```

You can enable or disable Windows authentication using either NTLM or Kerberos (as appropriate) by typing:

```
winrm set winrm/config/client @{TrustedHosts="<local>"}
```

or

```
winrm set winrm/config/client @{TrustedHosts=""}
```

In addition to managing WinRM at the command line, you can manage the service by using Group Policy. Keep in mind that Group Policy settings might override any other settings you enter.

Chapter 17. Managing Exchange Organizations

Although components of early Exchange releases were split into different server roles for scaling out Exchange organizations, later releases of Exchange Server streamlined the server roles and architecture incrementally while still allowing you to fully scale Exchange organizations to meet the needs of enterprises of all sizes. The end result of this streamlining is that Mailbox servers now perform all role services, except for Edge Transport services. Thus, not only do Mailbox servers store the active database copy for a mailbox performs all the data processing, rendering, and transformation required, they also provide authentication, redirection, and proxy services as needed.

For connections, the supported protocols include HTTP, POP, IMAP, MAPI over HTTP, RPC over HTTP, and SMTP, but no longer include RPC. Exchange 2016 is designed to work with Microsoft Outlook 2010 and higher and also continues to support the Outlook Web App for mobile access. Rather than connecting to servers by using Fully Qualified Domain Names (FQDN) as was done in the past, Outlook 2010 and higher use Autodiscover to create connection points based on the domain portion of the user's primary SMTP address and the GUID of a user's mailbox.

Navigating Exchange 2016 Organizations

An *organization* is the root of an Exchange environment. It's the starting point for the Exchange hierarchy, and its boundaries define the boundaries of any Exchange environment.

Organizational Architecture

When you install Exchange Server 2016, you install your Exchange servers within the organizational context of the domain in which the server is a member. The physical site boundaries and subnets defined for Active Directory Domain Services are the same as those used by Exchange Server 2016, and the site details are determined by the IP address assigned to the server. If you are installing the first Exchange server in a domain, you set the name of the Exchange organization for that domain. The next Exchange server you install in the domain joins the existing Exchange organization automatically.

Exchange 2016 organizations natively have only two server types: Mailbox servers on the internal network and Edge Transport servers on a perimeter network. Typically, client connections are load balanced with your Mailbox servers acting as the front end for Exchange services, and also acting as the back end servicing the connections (see Figure 17-1).




FIGURE 17-1 A basic Exchange organization with load balancers.

Clients connect to Exchange using several different techniques (see Figure 17-2). All Outlook connections take place using either MAPI over HTTP or Outlook Anywhere (RPC over HTTP). Using these technologies simplifies the protocol stack and required namespaces. POP3/IMAP and Unified Messaging connections are handled by dedicated services, as are SMTP connections, which are handled by the Transport stack. All other client connections, such as those for Outlook Web App (OWA), are managed via the IIS instance running on your Mailbox servers and then directed as appropriate.



FIGURE 17-2 Connections points into Mailbox servers.

Front End Transport

Mail transport is provided by the Front End Transport service, which provides mailbox locator services and proxies incoming and outgoing SMTP messages, as shown in Figure 17-3.

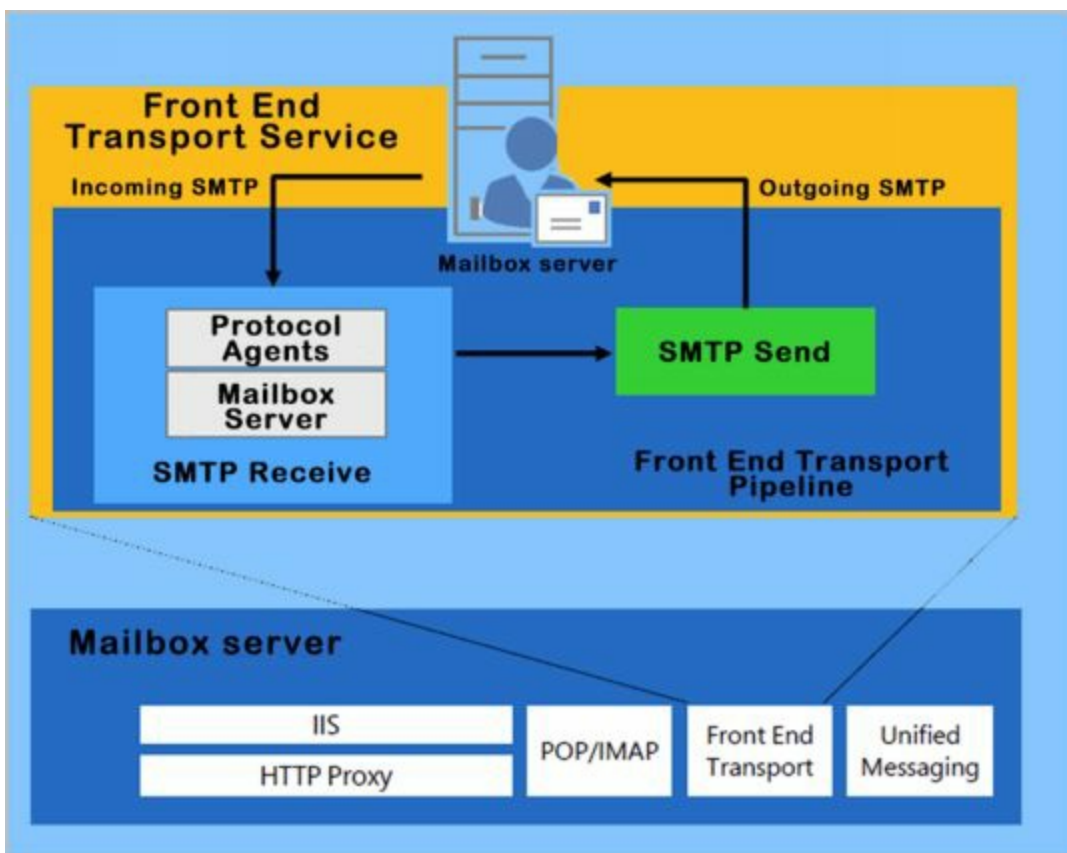


FIGURE 17-3 Front End Transport architecture

The Front End Transport service loads routing tables based on information from Active Directory and uses this information to route messages to the Transport service on Mailbox servers. The Mailbox server is selected based on the location of mailbox databases associated with the recipients.

A recipient is an entity that can receive Exchange mail and includes users, contacts, distribution groups, public folders, and resources (such as rooms and equipment used for scheduling). You refer to recipients as either *mailbox-enabled* or *mail-enabled*. Mailbox-enabled recipients (users and resources) have mailboxes for sending and receiving email messages. Mail-enabled recipients (contacts, distribution groups, and public folders) have email addresses but no mailboxes, which allow users in your

organization to send messages to mail-enabled recipients. Keep in mind that when you mail-enable a public folder and grant a user Send As permission on the folder, the user can send mail on behalf of the public folder.

In addition to users, contacts, groups, resources, and public folders, Exchange Server 2016 has two unique types of recipients: linked mailboxes and dynamic distribution groups. Basically, a linked mailbox represents a mailbox that is accessed by a user in a separate, trusted forest. A dynamic distribution group is a type of distribution group that you can use to build a list of recipients whenever mail addressed to the group is received, rather than having a fixed member list.

To manage recipients in your organization, you need to know these key concepts:

- **How address lists are used** Address lists are used to organize recipients and resources, making it easier to find the ones that you want to use, along with their related information. During setup, Exchange creates a number of default address lists, the most common of which is the global address list, which includes all the recipients in the organization. You can create custom address lists as well.
- **How email policies are used** Email address policies define the technique Exchange uses to create email addresses for users, resources, contacts, and mail-enabled groups. For example, you can set a policy that creates email addresses by combining an email alias with `@tvpres.com`. Thus, during setup of an account for William Stanek, the email alias *williams* is combined with `@tvpres.com` to create the email address *williams@tvpres.com*.
- **How retention policies are used** Retention policies are used to specify how long mail items remain in mailboxes and the actions to be taken when mail items reach their specified retention age. During setup, Exchange creates a default retention policy and this policy is applied automatically when you create an in-place archive mailbox for a user, providing no other retention policy is already applied.

The Routing tables used by the Front End Transport service contain a special list of Mailbox servers in the local Active Directory site. This list is based on the mailbox databases of message recipients. Routing in the front-end revolves around resolving message recipients to mailbox databases. For each mailbox database, the Front End Transport services looks up the routing destination.

Each routing destination has a delivery group, which is generally either a routable Database Availability Group (DAG), a Mailbox delivery group, or an Active Directory site, but can also be a group of connector source servers or a list of expansion servers for dynamic distribution groups. A Mailbox delivery group is a collection of one or more transport servers that are responsible for delivering messages to that routing destination. When the routing destination is a Mailbox delivery group, the delivery group may contain Exchange 2016 Mailbox servers, Exchange 2013 Mailbox servers, or Exchange 2010 Hub Transport servers.

The process by which the message is routed depends on the relationship between the source transport server and the destination delivery group. If the source transport server

is in the destination delivery group, the routing destination itself is the next hop for the message. The message is delivered by the source transport server to the mailbox database or connector on a transport server in the delivery group.

On the other hand, if the source transport server is outside the destination delivery group, the message is relayed along the least-cost routing path to the destination delivery group. In a complex Exchange organization, a message may be relayed either to other transport servers along the least-cost routing path, or directly to a transport server in the destination delivery group.

For an incoming message, the Front End Transport service selects a single Mailbox server to receive the message regardless of the number or type of recipients. If the message has a single recipient, a Mailbox server in the target delivery group is selected, with a preference based on the proximity of the Active Directory site. If the message has multiple recipients, the Front End Transport service uses the first 20 recipients to select a Mailbox server in the closest delivery group. If the message has no mailbox recipients, such as when the message is addressed to a distribution group, a Mailbox server in the local Active Directory site is randomly selected.

Back End Transport

The Back End Transport service runs on all Mailbox servers and is responsible for all mail flow within an Exchange organization, as shown in Figure 17-4.

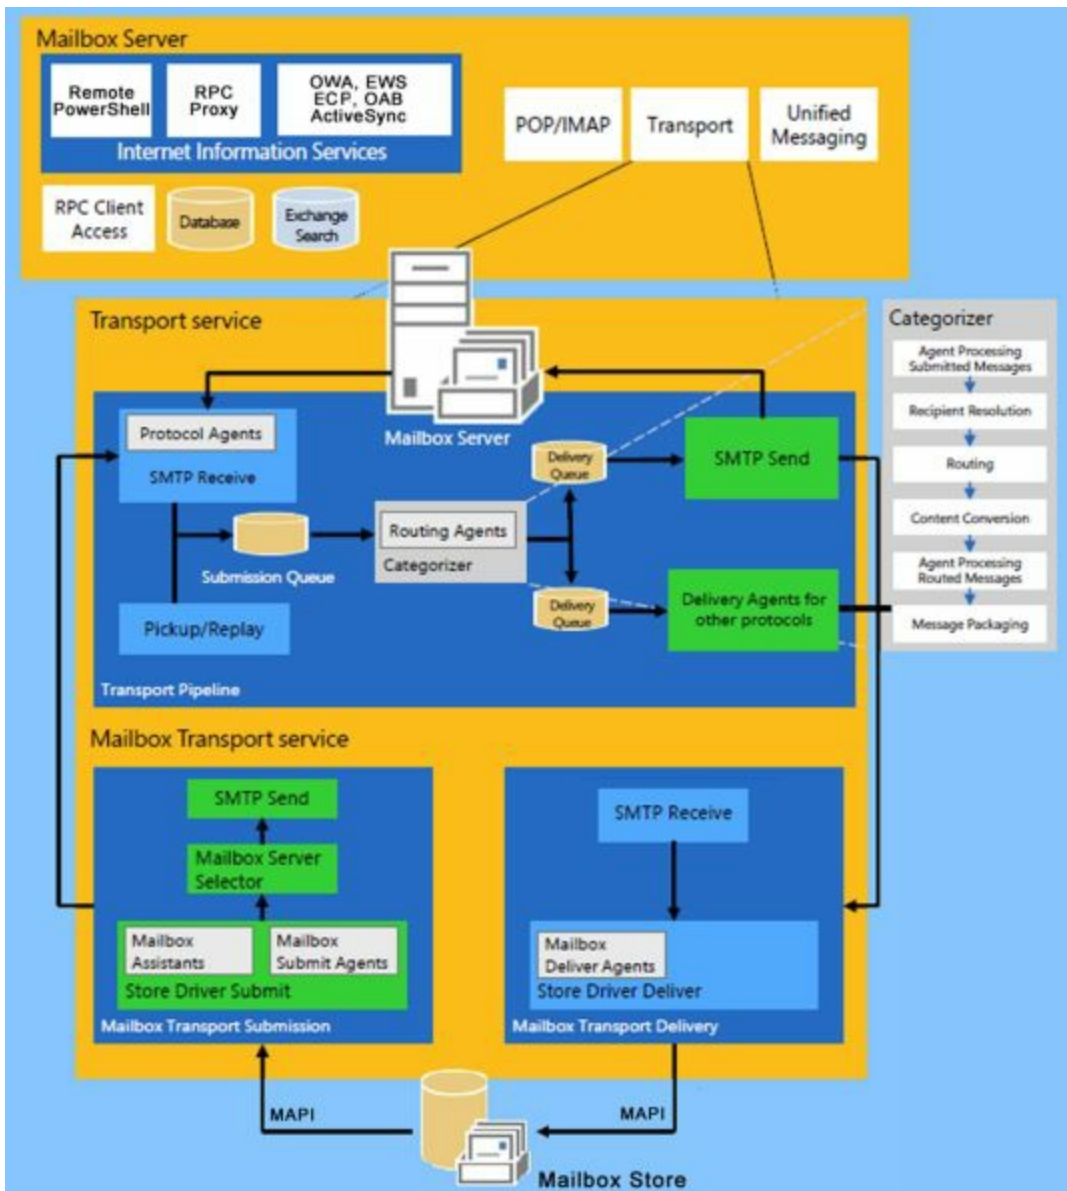


FIGURE 17-4 Back End Transport architecture

The Transport service relies on the Mailbox Transport service, which consists of two separate helper services: the Mailbox Transport Delivery service used with incoming messages and the Mailbox Transport Submission service used with outgoing messages. The transport service receives SMTP messages from the Transport service and establishes an RPC MAPI connection with the local mailbox database to deliver a message. The delivery service connects to the local mailbox database using RPC MAPI to retrieve messages and submits messages over SMTP to the Transport service.

The Mailbox Transport service only communicates with the Transport service and local mailbox databases. When this service receives a message for delivery it accepts the message if the recipient resides in an active copy of a local mailbox database. Otherwise, the service rejects the message and returns a non-delivery response to the Transport service for retrying delivery, generating a non-delivery report or rerouting the message.

When the Mailbox Transport service receives a message for submission, the service resolves the message recipients to mailbox databases. For each mailbox database, the service looks up the routing destination. Each routing destination has a delivery group,

which is either a routable DAG, a Mailbox delivery group, or an Active Directory site-- and the rest of the process continues as with the Front End Transport service.

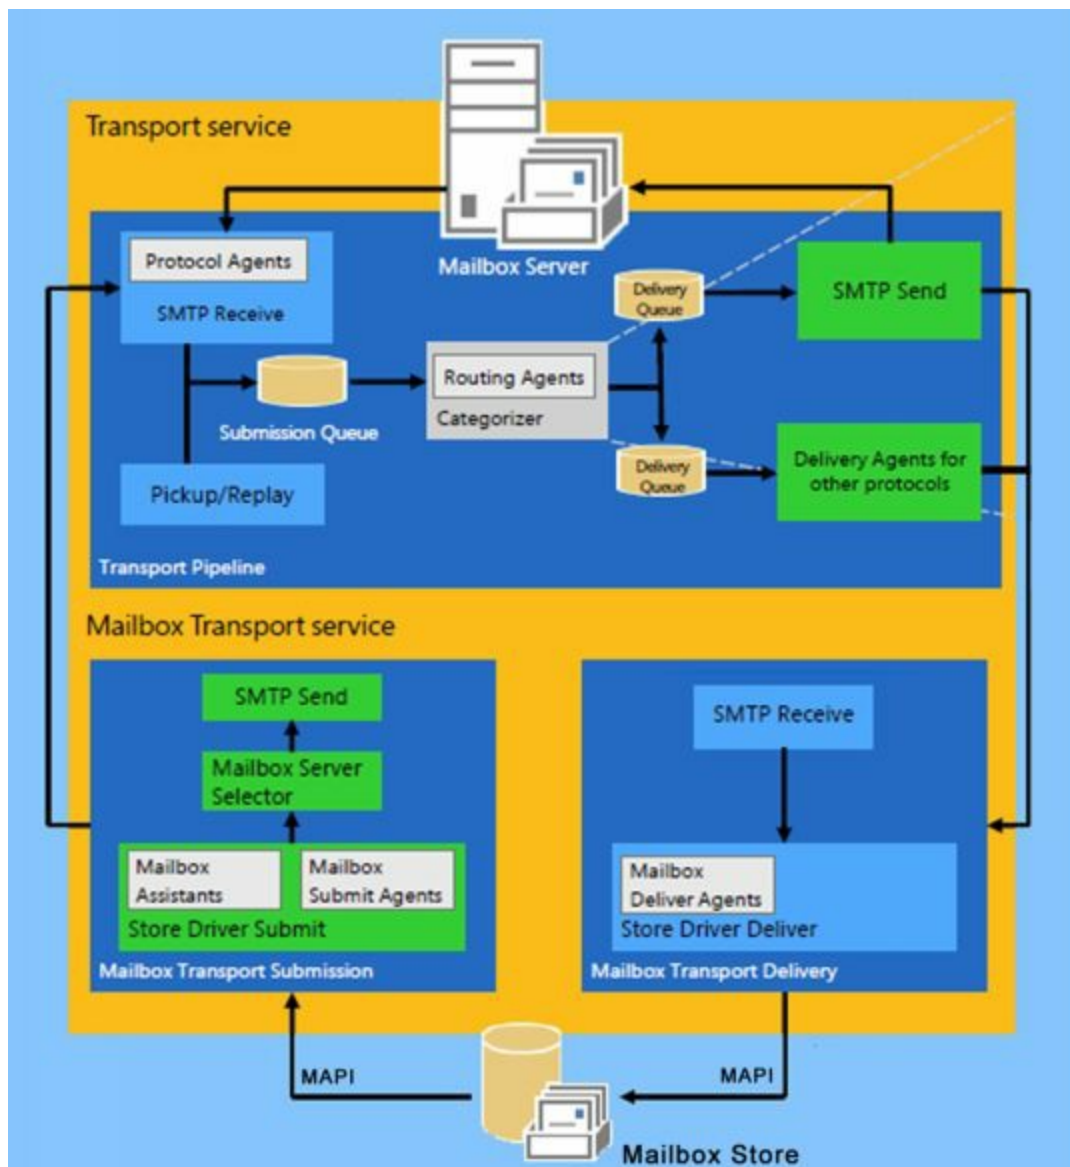


FIGURE 17-5 The Exchange Transport services

Exchange 2016 uses directory-based recipient resolution for all messages that are sent from and received by users throughout an Exchange organization. The Exchange component responsible for recipient resolution is the Categorizer. The Categorizer processes all email messages and uses the final recipient to determine what journaling policies, Information Rights Management policies, data loss prevention rules, and transport rules should be applied (see Figure 17-6).

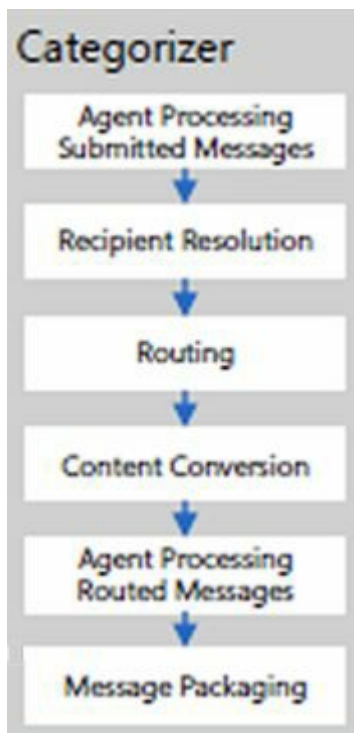


FIGURE 17-6 The Exchange Categorizer

The Categorizer must be able to associate every recipient in every message with a corresponding recipient object in Active Directory. All senders and recipients must have a primary SMTP address. If the Categorizer discovers a recipient without a primary SMTP address, it will determine what the primary SMTP address should be or replace a non-SMTP address. Replacing a non-SMTP address involves encapsulating the address in a primary SMTP address that will be used while transporting the message.

Understanding Exchange Routing

For routing messages, Exchange Server 2016 uses either Active Directory site-based routing or routing based on Database Availability Group (DAG) membership. The use of these routing approaches substantially changes the way you configure and manage Exchange Server 2016.

With Exchange Server 2016, site-based routing is possible because Exchange servers can determine their own Active Directory site membership and that of other servers by querying Active Directory. Using Active Directory for routing eliminates the need for Exchange to have its own routing infrastructure.

Routing Boundaries

Active Directory sites and DAGs are delivery group boundaries. When Mailbox servers aren't part of a DAG, they use site membership information to determine whether other Mailbox servers are located in the same site. This allows the Mailbox server to submit messages for routing and transport to another Mailbox server that has the same site membership. Site-based routing is also used for interoperability with Exchange 2013 and Exchange 2010.

When the destination delivery group is a collection of Mailbox servers in a single Active Directory site, the mailbox databases on those servers are the routing destinations. After a message is routed to the Transport service on a Mailbox server in the site, the Transport service routes the message to the Mailbox Transport service on the Mailbox server in the site that has the active copy of the destination mailbox database. The Mailbox Transport service in this server then delivers the message to the local mailbox database.

As routing destinations and delivery groups are separated by the major release version of Exchange, the Active Directory site may contain multiple Mailbox delivery groups. Specifically, each major release version of Exchange deployed in a particular site will have one delivery group. Regarding routing and delivery, keep the following in mind:

- Mailbox databases on Exchange 2010 Mailbox servers are serviced by Exchange 2010 Hub Transport servers in the site. After a message is routed to a random Hub Transport server in the site, the store driver on that server uses RPC to deliver the message into the mailbox database.
- Mailbox databases on Exchange 2013 Mailbox servers are serviced by the Transport service on Exchange 2013 Mailbox servers in the site. After a message is routed to the destination Mailbox server in the site, the Transport service uses SMTP to transfer the message to the Mailbox Transport service, which then uses RPC to deliver the message into the local mailbox database.
- Mailbox databases on Exchange 2016 Mailbox servers are serviced by the Transport service on Exchange 2016 Mailbox servers in the site. After a message is routed to the destination Mailbox server in the site, the Transport service uses SMTP to transfer the

message to the Mailbox Transport service, which then uses RPC to deliver the message into the local mailbox database.

When the destination delivery group is a routable DAG, the mailbox databases in the DAG are the routing destinations. After a message is routed to the Transport service on a Mailbox server in the DAG, the Transport service routes the message to the Mailbox Transport service that has the active copy of the destination mailbox database. The Mailbox Transport service in this server then delivers the message to the local mailbox database. As the DAG itself is the delivery group boundary rather than the Active Directory site associated with a particular Mailbox server, Mailbox servers may be physically located in more than one site even though they are members of the same delivery group.

IP Site Links

Exchange servers determine site membership by matching their assigned IP address to a subnet that is defined in Active Directory Sites and Services and associated with an Active Directory site. The Exchange server then uses this information to determine which domain controllers, Global Catalog servers, and other Exchange servers exist in that site, and it communicates with those directory servers for authentication, authorization, and messaging purposes. Exchange 2016 always tries to retrieve information about recipients from directory servers that are in the same site as the Exchange 2016 server.

TIP In Active Directory, you can associate a site with one or more IP subnets. Each subnet that is part of a site should be connected over reliable, high-speed links. You should configure any business locations connected over slow or unreliable links as part of separate sites. Because of this, individual sites typically represent well-connected local area networks (LANs) within an organization, and wide area network (WAN) links between business locations typically mark the boundaries of these sites. Sites cannot have overlapping subnet configurations because replication and message routing would not work correctly.

As Figure 17-7 shows, Active Directory sites are connected through IP site links, which can connect two or more sites. Each site link has a specific schedule, interval, and cost. The schedule and interval determine the frequency of Active Directory replication, and the cost value determines the cost of using the link relative to other links that might be available. Active Directory replication uses the link with the lowest cost when multiple paths to a destination exist. The cost of a route is determined by adding together the cost of all site links in a transmission path. Administrators assign the cost value to a link based on relative network speed, available bandwidth, and reliability compared to other available connections. By default, IP site links always allow traffic to flow into or out of a site.

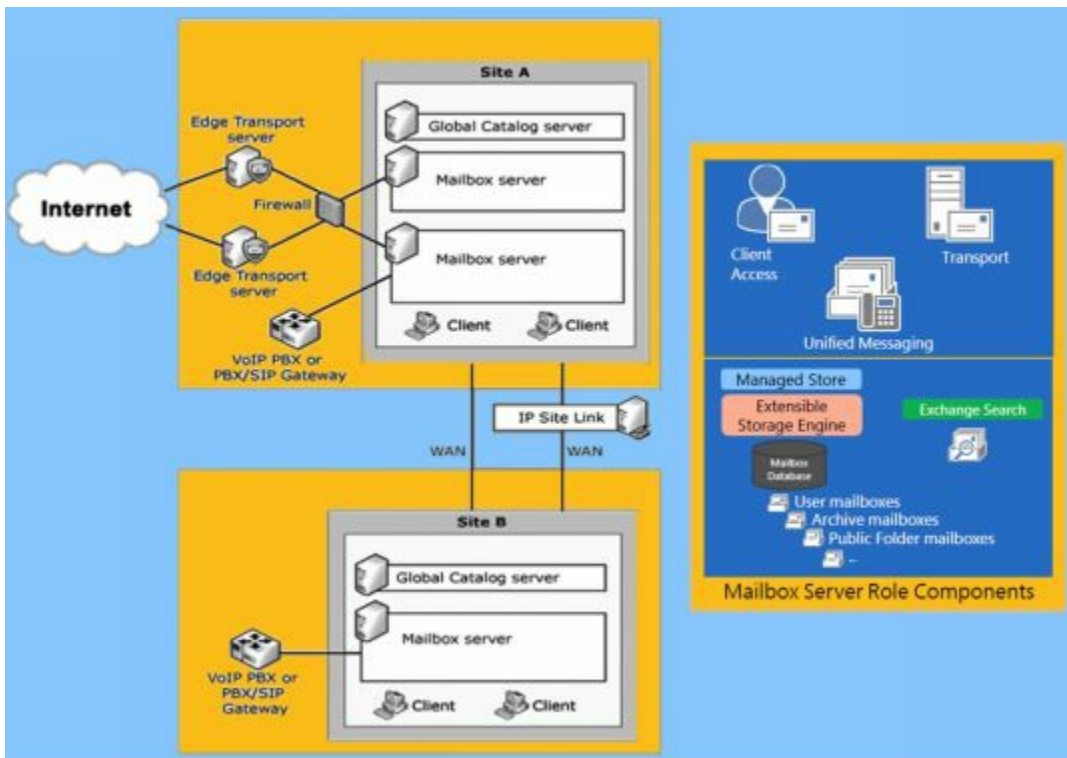


FIGURE 17-7 Message traffic between sites is routed over IP site links.

In large enterprises, message traffic might have to travel through multiple sites to get from the source site to a destination site. When transferring messages from one site to another through other sites, a transport server always tries to connect directly to a transport server in the destination site; therefore, messages are not relayed through each transport server in each site in the link path. Instead, the messages go directly from the transport server in the originating site across the link to the transport server in the destination site.

If the originating server cannot connect directly to a transport server in the destination site, the originating transport server uses the link cost to determine the closest site at which to queue the message. Messages queue until they are processed by the transport server and relayed to their destination. When Edge Transport servers are subscribed to an Active Directory site, the subscribed Edge Transport servers aren't accessible from other Active Directory sites.

The transport server can also use the site link information to optimize the routing of messages that users send to multiple recipients. The transport server expands a distribution list and creates multiple copies of a message only when multiple paths are in the routing topology. This feature is called *delayed fan-out*.

Delayed fan-out is used only when the delivery group is an Active Directory site. When multiple recipients share part of the routing path, delayed fan-out tries to reduce the number of message copies, thereby reducing the number of message transmissions.

Cross-Premises Routing

Exchange Online is a cloud service, meaning the service is provided via the Internet, and it allows you to outsource all or part of your Exchange services. Exchange Online

differs from Exchange on-premises (the standard implementation) in several fundamental ways. The Exchange Online hardware resides elsewhere and users access their mailboxes over the Internet; however, administrators still retain control and management over the outsourced mailboxes.

You can simultaneously connect to and manage both online and on-premises configurations in the Exchange Admin Center. Exchange Online has some advantages over an Exchange on-premises implementation, but has disadvantages as well. For users, Exchange Online provides:

- [Mailbox hosting](#)
- [ActiveSync](#)
- [Microsoft Outlook Anywhere](#)
- [Microsoft Outlook Web App](#)
- [Spam filtering](#)

For administrators, Exchange Online provides:

- [Service Level Agreements](#)
- [Storage quotas](#)
- [Automatic backups](#)
- [Automatic archiving](#)

It's important to point out that what Exchange Online doesn't provide is immediacy of access. Users must always be connected to the Internet to get their mail. Messages typically are routed and transferred across the Internet, which can cause delays. Exchange Online also does not offer Exchange voice mail as well as some other popular features.

As you work to configure your Exchange organization, it's important to keep in mind that Exchange Online is not an all-or-nothing implementation. You can host some mailboxes online and others on-premises. Exchange Server 2016 makes it easy to manage mailboxes regardless of where they are located. Before you transition mailboxes off-site, however, you'll probably want to perform a trial with a limited subset of users while keeping mailboxes for executives and most managers in house. In fact, it's a good idea to always keep mailboxes for executives and other high-level managers in house.

Exchange Server 2016 uses cross-premises routing to transfer messages between on-premises and hosted mailboxes. If you send a message to a user with a hosted mailbox, your organization's transport servers will route the message across the Internet to the hosted Exchange server. If you send a message to a user with an on-premises mailbox, your organization's transport servers will route the message across your organization to the appropriate Exchange server.

Exchange provides features for migrating mailboxes from online to on-premises environments and vice versa. During the migration, a mailbox might temporarily exist in both locations but when Exchange completes the migration, the mailbox exists only in the destination environment. Outlook 2010 and higher include an Autodiscover feature

that automatically connects messaging clients to the correct Exchange server. This feature utilizes the user's SMTP email address during automatic discovery to determine where the mailbox is currently located.

Normally, Autodiscover works very well; however, a conflict could occur if a user has a mailbox both in Exchange Online and in Exchange on-premises or if a user has the same primary SMTP email address in Exchange Online and Exchange on-premises. In these scenarios, the Autodiscover feature normally does not configure Outlook for the Exchange Online environment and instead uses Exchange on-premises, which has priority over Exchange Online when there is a conflict and the user's computer is connected to the Active Directory domain. To resolve the problem, delete the original mailbox from its first location as soon as possible after a mailbox migration. If a user needs both an online and on-premises mailbox, do not use the same primary SMTP email address for both Exchange Online and Exchange on-premises.

Understanding Data Storage in Exchange Server 2016

Exchange Server stores information in several locations, including:

- Active Directory data store
- Exchange Server store
- Exchange Server queues

Working with the Active Directory Data Store

The Active Directory data store contains most directory information for Exchange configurations and recipients as well as other important directory resources. Domain controllers maintain the data store in a file called Ntds.dit. The location of this file is set when Active Directory is installed and should be on an NTFS file system drive formatted for use with Windows Server. Domain controllers save some directory data separately from the main data store. Two key concepts on which to focus when looking at Active Directory are multimaster replication and Global Catalog servers.

Using Multimaster Replication

Domain controllers replicate most changes to the data store by using multimaster replication, which allows any domain controller to process directory changes and replicate those changes to other domain controllers. Replication is handled automatically for key data types, including the following:

- **Configuration data** Describes the topology of the directory, and includes a list of important domain information
- **Domain data** Contains information about objects within a domain, such as users, groups, and contacts
- **Schema data** Describes all objects and data types that can be stored in the data store

Using Global Catalogs

Active Directory information is also made available through global catalogs. Global catalogs are used for information searches and, in some cases, domain logon. A domain controller designated as a Global Catalog server stores a full replica of all objects in the data store (for its host domain).

The first domain controller installed in a domain is designated as the Global Catalog server by default. Consequently, if only one domain controller is in the domain, the domain controller and the global catalog are on the same server; otherwise, the global catalog is on the domain controller configured as such.

Global catalogs are primarily used for information searches. Searches in the global

catalog are efficient and can resolve most queries locally, thus reducing the network load and allowing for quicker responses. With Exchange, the global catalog can be used to execute Lightweight Directory Access Protocol (LDAP) queries for dynamic distribution groups. The members of the distribution group are based on the results of the query sent to the Global Catalog server rather than being fixed.

Why use LDAP queries instead of a fixed member list? The idea is to reduce administrative overhead by being able to dynamically determine the members of a distribution group. Query-based distribution is most efficient when the member list is relatively small (fewer than 100). If the member list has potentially hundreds or thousands of members, however, dynamic distribution can be inefficient and might require a great deal of processing to complete.

At a high-level, dynamic distribution works like this:

1. When email messages that are addressed to the group are received, the Exchange Categorizer (a transport component) sends the predefined LDAP query to the Global Catalog server for the domain.
2. The Global Catalog server executes the query and returns the resulting address set.
3. The Exchange Categorizer then uses the address set to generate the recipient list and deliver the message. If the Categorizer is unable to generate the list for any reason—for instance, if the list is incomplete or an error was returned—the Categorizer might start the process over from the beginning.

Using Dedicated Expansion Servers

To make the dynamic query process more efficient, Exchange 2016 shifts the processing requirements from Global Catalog servers to dedicated expansion servers by specifying a collection of one or more expansion servers as a delivery group. Unlike Mailbox delivery groups, this special delivery group can contain a mix of Exchange 2016 Mailbox servers, Exchange 2013 Mailbox servers and Exchange 2010 Hub Transport servers.

The routing destination is still the ultimate destination for a message. A distribution group expansion server is the routing destination when a distribution group has a designated expansion server that's responsible for expanding the membership list of the group. As with other types of routing, how the message is routed depends on the relationship between the source transport server and the destination delivery group. Keep in mind that when a distribution group expansion server is the routing destination, the distribution group is already expanded when a message reaches the routing stage of categorization on the distribution group expansion server. Therefore, the routing destination from the distribution group expansion server is always a mailbox database or a connector.

By default, Exchange 2016 uses the closest Exchange server that has the Mailbox server role installed as the dedicated expansion server. Because routing destinations and

delivery groups can also include Exchange 2013 Mailbox servers and Exchange 2010 Hub Transport servers in mixed environments, Exchange 2013 and Exchange 2010 servers could perform distribution group expansion in mixed Exchange organizations.

On occasion, you might want to explicitly specify the dedicated expansion server to handle expansion processing for some or all of your dynamic distribution groups in order to manage where the related processing occurs thereby shift the processing overhead from other servers to this specified server. You can specify a dedicated expansion server for a dynamic distribution group using the `-ExpansionServer` Parameter of the `Set-DynamicDistributionGroup` cmdlet.

Navigating the Exchange Information Store

The Exchange store is the core storage repository for managing Exchange databases. Unlike previous releases of Exchange, Exchange 2016 has only one type of database: the mailbox database. Mailbox databases contain the data, data definitions, indexes, flags, checksums, and other information that comprise mailboxes in your Exchange organization.

Data Storage Components

Exchange 2016 supports many types of mailboxes, including:

- **Arbitration mailbox** An arbitration mailbox is used to manage approval requests, such as handling moderated recipients and distribution group membership approval.
- **Archive mailbox** An alternative mailbox used to store historical mail items.
- **Discovery mailbox** A resource mailbox that is the target for Discovery searches.
- **Equipment mailbox** A resource mailbox for equipment scheduling.
- **Forwarding mailbox** A mailbox that can receive mail and forward it off-site.
- **Linked mailbox** A mailbox for a user from a separate, trusted forest.
- **Public folder mailbox** A shared mailbox for storing public folder data.
- **Room mailbox** A resource mailbox for room scheduling.
- **Shared mailbox** An alternative mailbox that is shared by multiple users, such as a general mailbox for customer inquiries.
- **User mailbox** The primary mailbox type for users to store mail items.

The Information Store (`Microsoft.Exchange.Store.Service.exe`) is written in C# and is fully integrated with the Microsoft Exchange Replication service (`MSEExchangeRepl.exe`). Officially, the store is referred to as the Managed Store. Although the Microsoft Exchange Information Store service hosts the Exchange store, which uses the Extensible Storage Engine (ESE) as the database engine, the management of the store is divided between the store service and the replication service. As you'd expect, the store service handles the primary store functions while the replication service provides replication and ancillary functions, including log shipping, log replay, log truncation and database seeding operations. The replication service also is responsible for all service availability for Mailbox servers.

The Active Manager component of the replication service is responsible for failure monitoring and failover within DAGs. The Active Manager is also responsible for message resubmissions from the shadow redundancy safety net. As examples, automatic resubmission of messages can occur after you activate the lagged copy of a mailbox database as well as after failover of a mailbox database in a DAG. Every Mailbox server runs Active Manager inside the replication service. If a Mailbox server isn't part of a DAG, the server has a single, Standalone Active Manager. In a DAG, there are two Active Manager roles: Primary Active Manager and Standby Active Manager. The Primary Active Manager determines which database copies are active and which are passive and also handles failover and notifies other members of topology changes.

The VSS writer in the replication service, named the Microsoft Exchange Writer, is responsible for backing up active and passive mailbox database copies and for restoring backed up database copies. Although this writer runs within the replication service, it is the store service that advertises the availability of the VSS writer. Thus, both the store service and the replication service must be running to back up and restore Exchange databases.

After a database backup, the transaction logs are usually truncated as the data is no longer needed for recovery; however, if backups aren't being taken, logs aren't truncated and you can prevent a buildup of logs by enabling circular logging for replicated databases. Exchange can use standard circular logging or continuous replication circular logging.

With standard circular logging, which is performed and managed by the store service, the Extensible Storage Engine (ESE) doesn't create additional log files because the current log file is overwritten when needed.

Combining standard circular logging with continuous replication is referred to as continuous replication circular logging. This type of logging is performed and managed by the replication service with a goal of maintaining log continuity. Logs are deleted only when they are no longer needed for replication.

The Managed Store

The Managed Store uses the worker process model. To isolate any issues with the Managed Store to a particular database, each database runs under its own process. Exchange Server uses transactions to control changes in databases and as with traditional databases, these transactions are recorded in a transaction log. Exchange Server then commits or rolls back changes based on the success of the transaction. The facility that manages transactions is the store service.

When working with databases, keep the following in mind:

- Each enterprise-level Mailbox server can have up to 100 databases (including both active and passive databases), with a maximum size per database of 64 terabytes (TB) — limited only by hardware.
- Each Mailbox server can be a member of only one database availability group and

can host only one copy (either the active or passive copy) of a particular database. Because each group can have up to 16 copies of a database, up to 16 different servers can be part of a database availability group.

To create a new mailbox database, you need about 50 megabytes (MBs) of free disk space. The files required by the database use a minimum of 23 MBs of disk space, and you'll need the extra space during creation and for read/write operations.

Other key concepts to focus on when working with the Exchange store and databases are the following:

- How Exchange server data files are used
- How data is stored in Exchange database files

Exchange Server Data Files

With Exchange Server 2016, Mailbox servers have a single database file for each mailbox database. Exchange 2016 stores all messages and attachments in the primary data file. Because attachments are encapsulated and written in binary format, you don't need to convert them to Exchange format. Exchange Server uses a link table within the database to reference the storage location of attachments within it.

As Figure 17-8 shows, each database has a primary data file and several other types of shared working files and transaction logs.

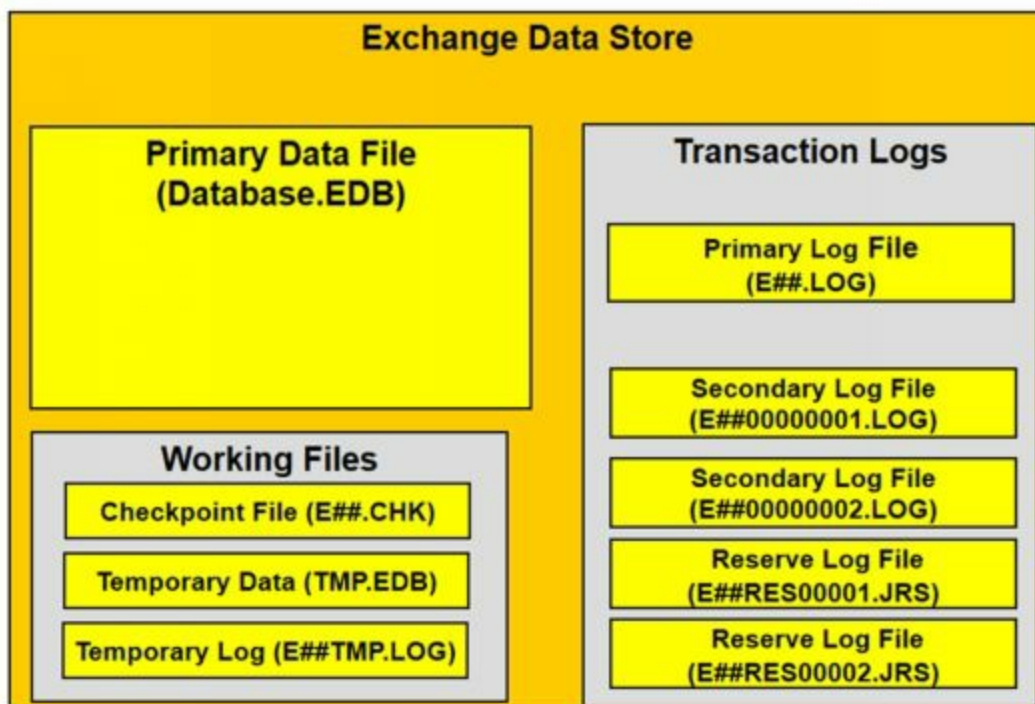


FIGURE 17-8 The Exchange data store has primary data files for each database as well as working files.

The file types are used as follows:

- **Primary data file (Database.edb)** A physical database file that holds the contents of the data store. By default, the name of the data file is the same as the name of the

associated data store with the .edb file extension added; however, you can rename a database without renaming the database file.

- **Checkpoint file (E##.chk)** A file that tracks the point up to which the transactions in the log file have been committed to databases in the storage group. Generally, the name of the checkpoint file is derived from the database prefix.
- **Temporary data (Tmp.edb)** A temporary workspace for processing transactions.
- **Current log file (E##.log)** A file that contains a record of all changes that have yet to be committed to the database. Generally, the name of the log file is derived from the database prefix.
- **Preprovisioned log file (E##tmp.log)** The next preprovisioned log, which is created in advance.
- **Secondary log files (E##00000001.log, E##00000002.log, ...)** Additional log files that are used as needed. Up to a billion unique log files can be created for each database.
- **Reserve log files (E##Res00001.jrs, E##Res00002.jrs, ...)** Files that are used to reserve space for additional log files if the current log file becomes full.
- **Temporary log (E##tmp.log)** A temporary workspace for logging.

By default, the primary data file, working files, and transaction logs are all stored in the same location. On a Mailbox server, you'll find these files in a per-database subfolder of the %SystemRoot%\Program Files\Microsoft\Exchange Server\V15\Mailbox folder. Although these are the main files used for the data store, Exchange Server uses other files, depending on the roles for which you have configured the server.

Data Storage in Exchange Databases

Exchange uses object-based storage. The primary data file contains several indexed tables, including a data table that contains a record for each object in the data store. Each referenced object can include object containers, such as mailboxes, and any other type of data that is stored in the data store.

Think of the data table as having rows and columns; the intersection of a row and a column is a field. The table's rows correspond to individual instances of an object, and the table's columns correspond to folders. The table's fields are populated only if a folder includes stored data. The data stored in fields can be a fixed length or a variable length.

Records in the data table are stored in data pages that have a fixed size of 32 kilobytes (KBs, or 32,768 bytes). The 32-KB page file size was changed from the 8-KB data pages used with Exchange Server 2007 to improve performance.

In an Exchange database, each data page has a page header, data rows, and free space that can contain row offsets. The page header uses the first 96 bytes of each page, leaving 32,672 bytes for data and row offsets. Row offsets indicate the logical order of rows on a page, which means that offset 0 refers to the first row in the index, offset 1 refers to the second row, and so on. If a row contains long, variable-length data, the data might not be stored with the rest of the data for that row. Instead, Exchange can store an

8-byte pointer to the actual data, which is stored in a collection of 32-KB pages that are written contiguously. In this way, an object and all its stored values can be much larger than 32 KB.

Changes to the mailbox database are written first to the transaction log and then committed to the database. The current active log file (E##.log) has a fixed size of 1 MB. When this log file fills up, Exchange closes the current active log file (E##.log) and renames it as E## NNNNNNNN .log (except when you are using circular logging). E##tmp.log is then renamed E##.log and becomes the current active log file.

The secondary log files are also limited to a fixed size of 1 MB. Exchange uses the reserve log files to reserve disk space for log files that it might need to create. Because several reserve files are already created, the transactional logging process is not delayed when additional logs are needed.

Exchange Server Message Queues

Exchange Server message queues are temporary holding locations for messages that are waiting to be processed. Two general types of queues are used:

- **Nonpersistent** Nonpersistent queues are available only when messages are waiting to be processed.
- **Persistent** Persistent queues are always available even if no messages are waiting to be processed.

With Exchange Server 2016, both Mailbox servers and Edge Transport servers store messages waiting to be processed in nonpersistent and persistent queues. The list that follows provides an overview of the queues used:

- A nonpersistent primary and shadow Safety Net / Transport dumpster queue on Mailbox and Edge Transports for each Active Directory site
- Nonpersistent *Delivery / Relay* queues found on Mailbox servers with one for each unique destination Mailbox server, connector, designated expansion server, non-SMTP gateway, etc.
- One nonpersistent Remote delivery queue on Edge Transport servers for each unique destination SMTP domain and smart host
- One nonpersistent Shadow redundancy queue on Mailbox and Edge Transport servers for each hop to which the server delivered the primary message
- One persistent Poison message queue found on Mailbox and Edge Transport servers
- One persistent Submission queue on Mailbox and Edge Transport servers
- One persistent Unreachable queue on Mailbox and Edge Transport servers

You can view top-level queues by clicking Queue Viewer in the Exchange Toolbox. You'll learn more about queues in Chapter 29, "Maintaining Exchange Server 2016."

When working with queues, Shadow redundancy and Safety Net are two important concepts that you need to understand. While shadow redundancy keeps a redundant copy of messages in transit, Safety Net keeps a redundant copy of a message after the message

is successfully processed. Thus, in effect, Safety Net takes over where shadow redundancy finishes.

Exchange Server 2016 implements shadow redundancy for queued messages. In the event of an outage or server failure, this feature works to prevent the loss of messages that are in transit by storing queued messages until the next transport server along the route reports a successful delivery of the message. If the next transport server doesn't report successful delivery, the message is resubmitted for delivery.

Shadow redundancy eliminates the reliance on the state of any specific Mailbox or edge server and eliminates the need for storage hardware redundancy for transport components. As long as redundant message paths exist in your routing topology, any transport component is replaceable and you don't have to worry about emptying a server's queues or losing messages due to transport failure.

In Exchange 2016, the Transport service makes a redundant copy of a message as soon as it receives it and then acknowledges receipt. In Exchange 2010, the Transport service would acknowledge receipt and then make a redundant copy of a message. Finally, it's important to note that it doesn't matter whether the sending server supports shadow redundancy. If Exchange 2016 determines that a message was lost in transit, Exchange delivers the messages using the redundant copy.

TIP Shadow redundancy uses less bandwidth than creating duplicate copies of messages on multiple servers. The only additional network traffic is the exchange of discard status between transport servers. Discard status indicates when a message is ready to be discarded from the transport database.

Exchange Server 2016 also implements Safety Net for queued messages. Safety Net replaces and enhances the transport dumpster available in Exchange 2010. By default, Safety Net stores copies of messages that were successfully processed by a Mailbox server for two days. For Mailbox servers that aren't part of DAGs, Safety Net stores copies of messages delivered to other Mailbox servers in the local Active Directory site. For Mailbox servers that are part of DAGs, Safety Net stores copies of messages delivered to other Mailbox servers in the DAG.

Since Safety Net uses shadow redundancy, it is always fully redundant. The Primary Safety Net queue stores the primary copy of a delivered message. The Shadow Safety Net queue stores a shadow copy of a delivery message. If the Primary Safety Net queue is unavailable for more than 12 hours, any messages that need to be redelivered are redelivered from the Shadow Safety Net queue.

When Mailbox servers are part of a DAG, Safety Net is used for some shadow redundancy functions. Previously, in a DAG, shadow redundancy would keep a copy of messages in the shadow queue until they were replicated to passive copies of the database. As Safety Net already has a copy of delivered messages, shadow redundancy doesn't need to keep another copy of these messages and messages can be resubmitted from Safety Net if necessary.

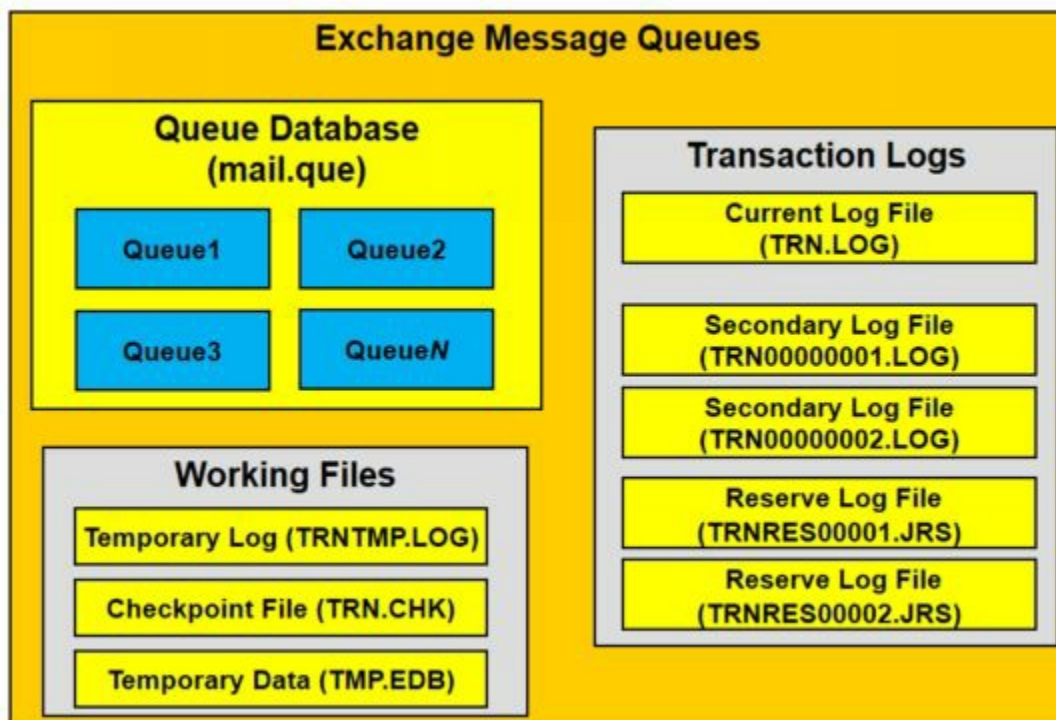


FIGURE 17-9 The Exchange message queues are all stored in a single database.

The various message queues are all stored in a single database (see Figure 17-9). Like the Exchange store, the message queues database uses the ESE for message storage as well as for data pages.

The database has a single data file associated with it and several other types of working files and transaction logs. These files are used as follows:

- **Primary data file (Mail.que)** A physical database file that holds the contents of all message queues.
- **Checkpoint file (Trn.chk)** A file that tracks the point up to which the transactions in the log file have been committed to the database.
- **Temporary data (Tmp.edb)** A temporary workspace for processing transactions.
- **Current log file (Trn.log)** A log file that contains a record of all changes that have yet to be committed to the database.
- **Preprovisioned log file (Trntmp.log)** The next preprovisioned log, which is created in advance.
- **Secondary log files (TRN00000001.log, TRN00000002.log, ...)** Additional log files that are used as needed.
- **Reserve log files (TRNRes00001.jrs, TRNRes00002.jrs, ...)** Files that are used to reserve space for additional log files if the current log file becomes full.

The facility that manages queuing transactions is the Microsoft Exchange Transport service (MSExchangeTransport.exe). Because logs used with message queues are not continuously replicated, these log files have a fixed size of 5 MB. Changes to the queue database are written first to the transaction log and then committed to the database. When the current active log (trn.log) file fills up, Exchange closes the file and renames it as TRN NNNNNNNN .log. Trntmp.log is then renamed Trn.log and becomes the current active log file.

Exchange uses the reserve log files to reserve disk space for log files that might need to be created. Because several reserve files are already created, this speeds up the transactional logging process when additional logs are needed.

By default, the data file, working files, and transaction logs are all stored in the same location. On a Mailbox server or Edge Transport server, you'll find these files in the %SystemRoot%\Program Files\Microsoft\Exchange Server\V15\TransportRoles\data\Queue folder.

Chapter 18. Implementing Availability Groups

Managing the information store is one of your most important tasks as a Exchange 2016 administrator. The Information Store (Microsoft.Exchange.Store.Service.exe) is written in C# and is fully integrated with the Microsoft Exchange Replication service (MSEExchangeRepl.exe) and the Microsoft Exchange DAG Management service (MSEExchangeDagMgmt.exe).

Every Mailbox server that is deployed in your organization has an information store, which can contain databases and information about database availability groups (DAGs). This chapter introduces databases and focuses on the management of database availability groups. After completing this chapter, you should know how to:

- [Enable, create, and use database availability groups](#)
- [Manage databases and their related transaction logs](#)
- [Improve Mailbox server availability](#)
- [Manage full-text indexing of Exchange databases](#)

See Chapter 19, “Configuring Exchange Databases,” to learn how to manage databases.

Building Blocks for High Availability

Microsoft built high availability and messaging resilience into the core architecture of Exchange 2016, providing a simple unified framework for high availability, management, and disaster recovery. This approach allows Exchange 2016 to improve continuous replication, provide a robust solution that doesn't require expensive clustering hardware, and reduce maintenance overhead.

The Extensible Storage Engine

Exchange Server 2016 uses Extensible Storage Engine (ESE) databases for mailbox storage. When you install a Mailbox server in an Exchange 2016 organization, this server's information store has a single, default mailbox database. Mailbox databases have a single, dedicated log stream, which is represented by a series of sequentially named log files. Each log file is 1 megabyte (MB) in size. In addition to log files, databases have several other types of files associated with them, including one or more checkpoint files, a temporary working file, and one or more transaction log files. Depending on the state of Exchange Server, you might see other working files as well.

NOTE Exchange 2016 does not use public folder databases. Public folders are now stored in a special type of mailbox.

When you create a mailbox database, you can specify separate folder locations to use for database files and transaction logs. Each database has content-indexing files associated with it as well. These files are generated by the Exchange Search service, which is enabled by default and running on all Mailbox servers. Exchange Search indexes new mail items in the transport pipeline or immediately after the items are created and delivered to a mailbox.

You use Exchange databases to ease the administrative burden that comes with managing large installations. For example, instead of having a single 10-terabyte (TB) database for the entire organization, you can create ten 1-TB databases that you can manage more easily.

TIP As a best practice, 2 TB is the largest recommended size for Exchange Server 2016 databases. Often you'll find that large databases make it easier to support the large mailboxes that might be required by your organization's managers and executives. Still, most mailboxes should be limited to between 2 GB and 10 GB in size.

When you create a mailbox database, you specify the name for the database, and this name sets the name of the primary database file as well. For example, if you create a mailbox database called EngineeringDept, the primary database file is set as EngineeringDept.edb. With Exchange Server 2016, the default location for database files is the same as the log folder. If you want a database to be in a different location, you can specify the location you want to use.

Separating database files and log files from the same database and putting them on different volumes backed by different physical disks can help you scale your organization while ensuring high performance and recoverability. When you are formatting volumes, be sure to consider using ReFS as opposed to NTFS as ReFS is more efficient and resilient than NTFS.

TIP Recoverability is a key reason for separating database files and log files. For example, in the case of a failure on a drive where a database is stored, the transaction logs needed for complete recovery would then be on a different (and probably functioning) drive. Whether you want to use this approach depends on the size and configuration of your Exchange Mailbox servers as well as the service level agreements with which you need to comply.

The many files associated with databases provide granular control over Exchange Server, and if you configure the data files properly, they can help you scale your Exchange organization efficiently while ensuring optimal performance. In a small implementation of Exchange, you might want to place all the data files on the same drive. As you scale from a small organization to a larger organization, you'll generally want to organize data according to databases, placing all the data for each database on physically separate drives. You can't always do this, however, in a small-to-medium sized organization with limited resources. For example, if you have ten 1-TB databases and only five data drives, you might want to have the five data drives configured as follows:

- Drive 1 with Database 1 and Database 2 and all related data files.
- Drive 2 with Database 3 and Database 4 and all related data files.
- Drive 3 with Database 5 and Database 6 and all related data files.
- Drive 4 with Database 7 and Database 8 and all related data files.
- Drive 5 with Database 9 and Database 10 and all related data files.

In a storage area network (SAN) implementation in which you are using logical unit numbers (LUNs) and don't know about the underlying disk structure, placing the databases on separate LUNs should be sufficient. To protect the data, you might want to consider using hardware RAID (redundant array of inexpensive disks), which is likely already implemented if you are using a SAN. However, if you configure a database availability group with multiple member servers that each have one or more copies of mailbox databases, you likely don't need to use any type of RAID, and you likely won't need daily backups either. Just remember that Microsoft recommends having at least three database copies in addition to the active copy.

REAL WORLD When you have multiple copies of your data on separate servers, you really might not need to create daily backups of your Exchange data. This doesn't mean that you won't need to create backups ever—it just means you might not need daily backups of Exchange data. You will probably still want to create regular backups of your Exchange servers and still create periodic full backups of all server and Exchange data to rotate to off-site storage as a safeguard

against catastrophe.

Database available groups can also make you rethink your use of SANs. Rather than having a single, massive (and likely very expensive) storage device, you might want to rely on a server's internal drives or multiple smaller (and likely much less complex) storage devices. One reason to use internal drives is that reliable, multiple-TB hard drives are becoming increasingly available, and several servers with multiple, large internal hard drives will likely cost a fraction of the price of a single massive SAN. If you use SANs, you might find that multiple smaller storage devices are better than a single, massive storage device because you'll then be protected against a single source of failure (the storage device) causing an outage on all your mailbox servers. I know, I know...the SAN should never go down, but it can (and does) happen.

The High Availability Framework

Exchange 2016 allows you to protect mailbox databases and the data they contain by configuring your mailbox databases for high availability automatically when you use database availability groups (DAGs). Database availability groups allow you to group databases logically according to the servers that host a set of databases.

Each Mailbox server can have multiple databases, and each database can have as many as 16 copies. A single database availability group can have up to 16 Mailbox servers that host databases and provide automatic database-level recovery from failures that affect individual databases. Any server in a database availability group can host a copy of a mailbox database from any other server in the database availability group.

Exchange 2016 integrates high availability and messaging resilience into the core architecture, providing a simple unified framework for both high availability and disaster recovery. This approach reduces the cost and complexity of deploying a highly available solution. How does this work? Exchange 2016 has enhanced continuous replication and has replaced clustering features in early releases of Exchange with a more robust solution that doesn't require expensive hardware and also requires less maintenance.

In early versions, Exchange was a clustered application that used the cluster resource management model for high availability. In contrast, Exchange 2016 is not a clustered application and therefore does not use the cluster resource model for high availability. Instead, Exchange 2016 uses its own internal high-availability model. Although some components of Windows Failover Clustering are still used, these components are now managed exclusively by Exchange 2016.

Early versions of Exchange supported continuous replication through several approaches, including Local Continuous Replication (LCR), Cluster Continuous Replication (CCR), and Standby Continuous Replication (SCR). LCR was a single-server solution for asynchronous log shipping, replay, and recovery. CCR combined the asynchronous log shipping, replay, and recovery features with the failover and

management features of the Cluster service, and it was designed for configurations in which you had clustered Mailbox servers with dedicated active and passive nodes. SCR was an extension of LCR and CCR that used the same log shipping, replay, and recovery features of LCR and CCR but was designed for configurations in which you used or enabled the use of standby recovery servers.

Some aspects of the continuous replication technology previously found in CCR and SCR are built into Exchange 2016, but the technology has changed substantially. Because storage groups have been removed from Exchange 2016, continuous replication operates at the database level. Exchange 2016 still uses an Extensible Storage Engine (ESE) database that produces transaction logs that are replicated and replayed into copies of mailbox databases. Because each mailbox database can have as many as 16 copies, you can have one or more database copies on up to 16 different servers.

When a Mailbox server is added to DAG, the server works with other members of the DAG to provide automatic recovery from failures that affect mailbox databases, including disk failures, server failures, and other critical failures. When a failure affecting a database occurs and a new database becomes the active copy automatically, this process is known as a failover. When an administrator establishes a database copy as the active mailbox database, this process is known as a switchover.

Failover and switchover occur at the database level for individual databases and at the server level for all active databases hosted by a server. When either a switchover or failover occurs, other Exchange 2016 server roles become aware of the switchover almost immediately and redirect client and messaging traffic automatically as appropriate.

For site resiliency, you use availability groups and deploy Mailbox servers in two or more Active Directory sites. If the sites are located in different datacenters with different network infrastructure, you can ensure mailbox data is protected from software, hardware and datacenter failures. When you use datacenter pairs, you'll want to deploy a single availability group with member servers located in each datacenter. For example, if you have two datacenters, you could deploy two servers in each datacenter and make each server a member of the same availability group. You would then have a single availability group with four member servers. To ensure full protection, each database would have four copies, with two copies in each datacenter. Of the four copies:

- 1 would be active and highly available
- 2 would be passive and highly available
- 1 would be passive and lagged

Lagging a database is an important concept to understand. When you lag a database copy, you configure Exchange to delay processing of log files for a designated database copy. Exchange then processes (or plays down) the log files with a preset time delay, such as 3 days. Because of the delay in processing, the log files of the lagged copy can be used to automatically recover from broken page references and other issues that

would otherwise require administrator intervention for recovery. For example, if Exchange detects that page patching is required, Exchange automatically replays the logs into the lagged copy to perform page patching.

With Exchange Server 2016, Replay Lag Manager is enabled by default and lagged database copy play down can be delayed based on disk latency to ensure active users aren't impacted. In most circumstances, you lag the database copy with the highest activation preference number to protect against catastrophic logical corruption of the database. As a best practice, you should use a replay lag time of at least 7 days and enable the Replay Lag Manager to provide dynamic log file play down of lagged copies when availability is compromised.

You can perform most management tasks for availability groups in the Exchange Admin Center. However, you have additional options when you work with Exchange Management Shell. Commands you can use to manage availability groups and their various features include:

- **Database availability group management**

- `Get-DatabaseAvailabilityGroup`
- `New-DatabaseAvailabilityGroup`
- `Remove-DatabaseAvailabilityGroup`
- `Set-DatabaseAvailabilityGroup`

- **Database copy management**

- `Add-MailboxDatabaseCopy`
- `Get-MailboxDatabaseCopyStatus`
- `Remove-MailboxDatabaseCopy`
- `Resume-MailboxDatabaseCopy`
- `Set-MailboxDatabaseCopy`
- `Suspend-MailboxDatabaseCopy`
- `Update-MailboxDatabaseCopy`

- **Database management**

- `Dismount-Database`
- `Get-MailboxDatabase`
- `Move-DatabasePath`
- `New-MailboxDatabase`
- `Remove-MailboxDatabase`
- `Set-MailboxDatabase`

- **Network configuration**

- `Get-DatabaseAvailabilityGroupNetwork`
- `New-DatabaseAvailabilityGroupNetwork`
- `Remove-DatabaseAvailabilityGroupNetwork`
- `Set-DatabaseAvailabilityGroupNetwork`

- **Switchover management**

- Move-ActiveMailboxDatabase

- Start-DatabaseAvailabilityGroup

- Stop-DatabaseAvailabilityGroup

- Restore-DatabaseAvailabilityGroup

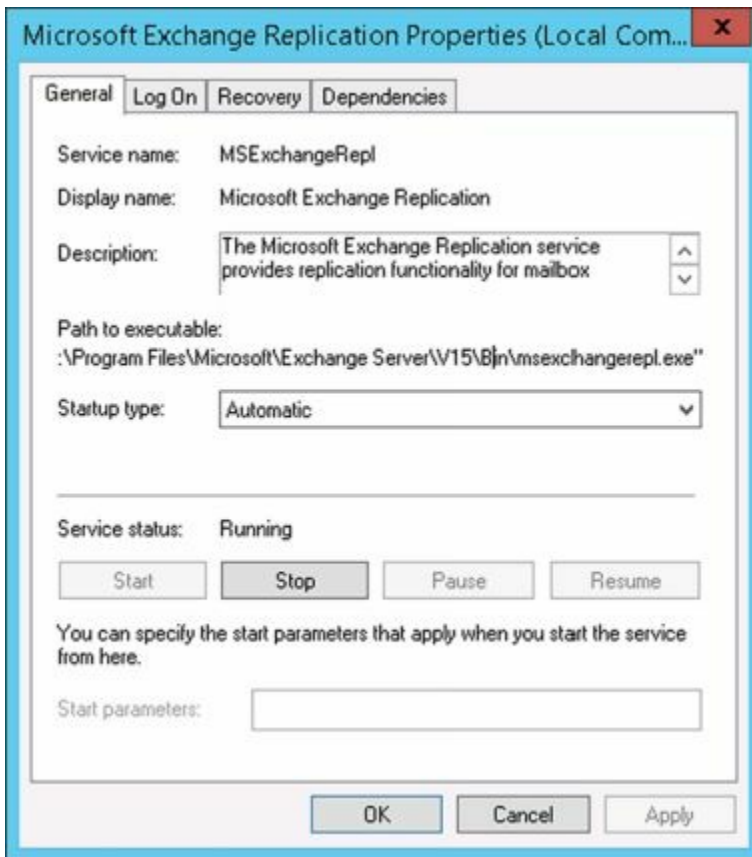
- **Server membership**

- Add-DatabaseAvailabilityGroupServer

- Remove-DatabaseAvailabilityGroupServer

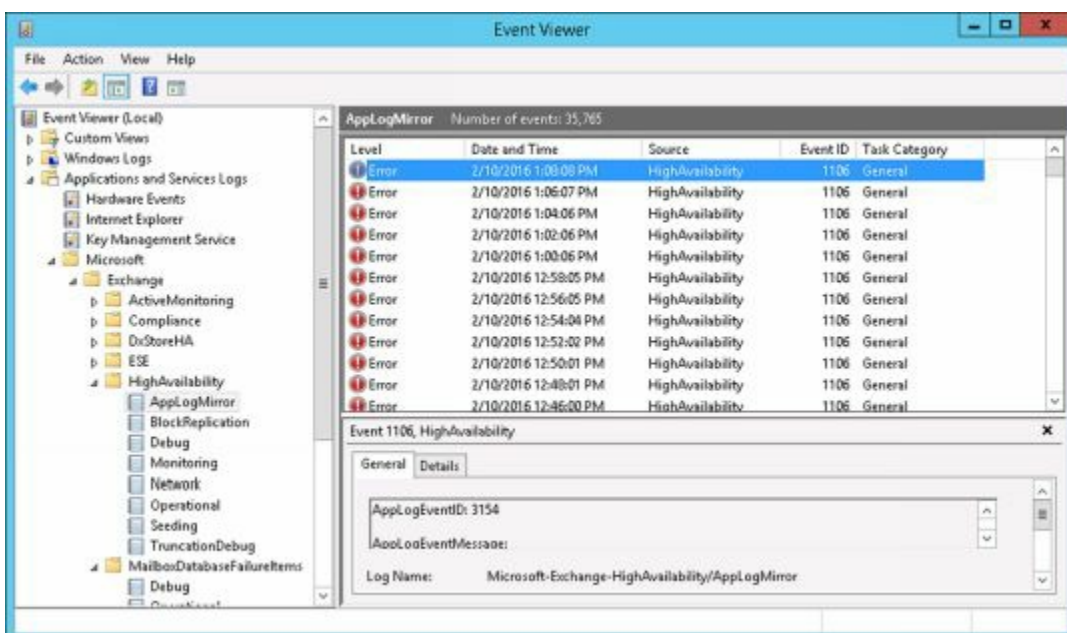
When planning database availability groups, keep in mind that you can create database copies only on Mailbox servers in the same database availability group that do not host the active copy of a database. An active copy differs from a passive copy in that it's in use and being accessed by users rather than offline. You cannot create two copies of the same database on the same server. Other guidelines to keep in mind when working with database copies include the following:

- All Mailbox servers in a database availability group must be in the same Active Directory domain. Database copies can be created 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).
- You cannot replicate a database outside a database availability group. This means Mailbox databases can be replicated only to other Mailbox servers in the same database availability group.
- 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.



The service responsible for replicating databases is the Microsoft Exchange Replication (MSExchangeRepl) service. The replication service and components that run within the service, including Active Manager, the TCP listener, and the Volume Shadow Copy Service writer, write results to the event logs.

In Event Viewer, you can find these logs by navigating to Applications and Services Logs > Microsoft > Exchange > High Availability. In these logs, you'll find details on database actions, such as database mount operations, log truncation, and cluster action within DAGs. Events related to failures that affect replicated mailbox databases are written to the logs under Applications and Services Logs > Microsoft > Exchange > MailboxDatabaseFailureItems.



Cluster Components

In Exchange 2016, Active Manager provides the resource model and failover management features previously provided by the Cluster service. When you create your first database availability group in an Exchange organization, Exchange creates a Windows Failover Cluster, but there are no cluster groups for Exchange and no storage resources in the cluster.


Failover Cluster Manager shows only basic information about any cluster, which includes the cluster name and networks, and the quorum configuration. Cluster nodes and networks will also exist, and their status can be checked in Failover Cluster Manager; however, Exchange manages all cluster resources, including nodes and networks.



REAL WORLD Failover Cluster Manager is the primary management tool for working with the Cluster service. Although you need to use the Exchange Management tools to view and manage database availability groups and related features, Failover Cluster Manager does show the status of clustering.

- By selecting the cluster name in the left pane, you get a quick overview of the cluster configuration, including the current quorum configuration, which can be either Node Majority or Node and File Share Majority depending on the number of nodes in the database availability group.
- By selecting the Nodes entry in the left pane, you can quickly check the status of all the nodes in the database availability group.
- By expanding the Networks entry in the left pane and then selecting available cluster networks, you can check the status of the network as well as individual network connections.
- By selecting the Cluster Events node, you can check the event logs on all cluster nodes for errors and warnings.

Exchange makes use of the cluster's node and network management functions. You can check the node and network status in Exchange Admin Center by selecting Servers > Database Availability Groups.



Active Manager Framework

Active Manager runs on all Mailbox servers as a subcomponent of the Microsoft Exchange Replication service. On Mailbox servers that aren't part of a DAG, Active Manager operates as a Standalone Active Manager. On Mailbox servers that are members of a DAG, Active Manager operates as either a primary role holder or a standby secondary role holder with respect to a particular database. The primary role holder, referred to as the Primary Active Manager, decides which database copies will be active and which copies to activate. It also receives topology change notifications and reacts to server failures. Only one copy of a database can be active at any given time, and that copy can be mounted or dismounted.

The group member that holds the primary role is always the member that currently owns the cluster quorum resource and the default cluster group. If the server that owns the cluster quorum resource fails, the primary role automatically moves to another server in the group and that server takes ownership of the default cluster group. Before you take the server that hosts the cluster quorum resource offline for maintenance or an upgrade, you must first move the primary role to another server in the group.

Secondary role holders, referred to as Standby Active Managers, provide information about which server hosts the active copy of a mailbox database to other Exchange components. The secondary role holder detects failures of replicated, local databases and the local information store, and it issues failure notifications to the primary role holder and asks the primary role holder to initiate a failover. The secondary role holder does not determine which server takes over, nor does it update the database location state with the primary role holder. With respect to its local system, the primary role holder also performs the functions of the secondary role by detecting local database and local information store failures and issuing related notifications.

Active Manager determines which database copy should be activated by attempting to locate a mailbox database that has characteristics similar to the following:

- The database has a status of Healthy, DisconnectedAndHealthy, DisconnectedAndResynchronizing, or SeedingSource. @techjob
- The database has a content index with a status of Healthy.
- The database has a copy queue length that is less than 10 log files.
- The database has a replay queue length of less than 50 log files.
- The server hosting the database has all components in a healthy state.

If no database copy meets all of these criteria, Active Manager continues looking for the best choice by lowering the selection requirements through successive iterations. Active Manager uses the managed availability framework to perform health checks. After one or more copies have been selected, Active Manager attempts to copy any missing log files from the original source to a potential new active copy by using a process called attempt copy last logs (ACLL). Once the ACLL process is complete, Active Manager compares the value of the AutoDatabaseMountDial property for Mailbox servers hosting copies of the database to the copy queue length of the database being activated. If the value of the AutoDatabaseMountDial property is greater than the number of missing log files, the Primary Active Manager tries to activate the next best copy (if one is available).

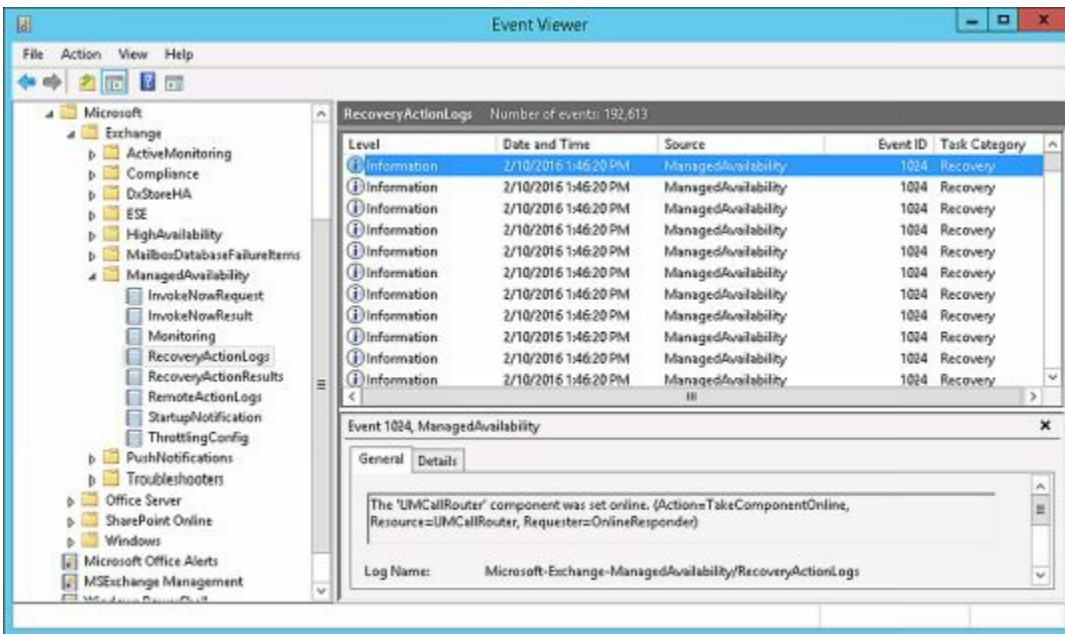
If the value of the AutoDatabaseMountDial property is equal to or less than the number of missing log files, the Primary Active Manager issues a mount request. At this point, either the database mounts and is made available to clients or the database doesn't mount and the Primary Active Manager tries to activate the next best copy (if one is available).

Managed Availability Components

In Exchange 2016, the active monitoring and high availability functions are integrated into a single architecture called managed availability, which is implemented on Mailbox servers. Managed availability is a framework that includes a probe engine for taking measurements and collecting data, a monitor engine for determining the status of Exchange components, and a responder engine for taking recovery actions.

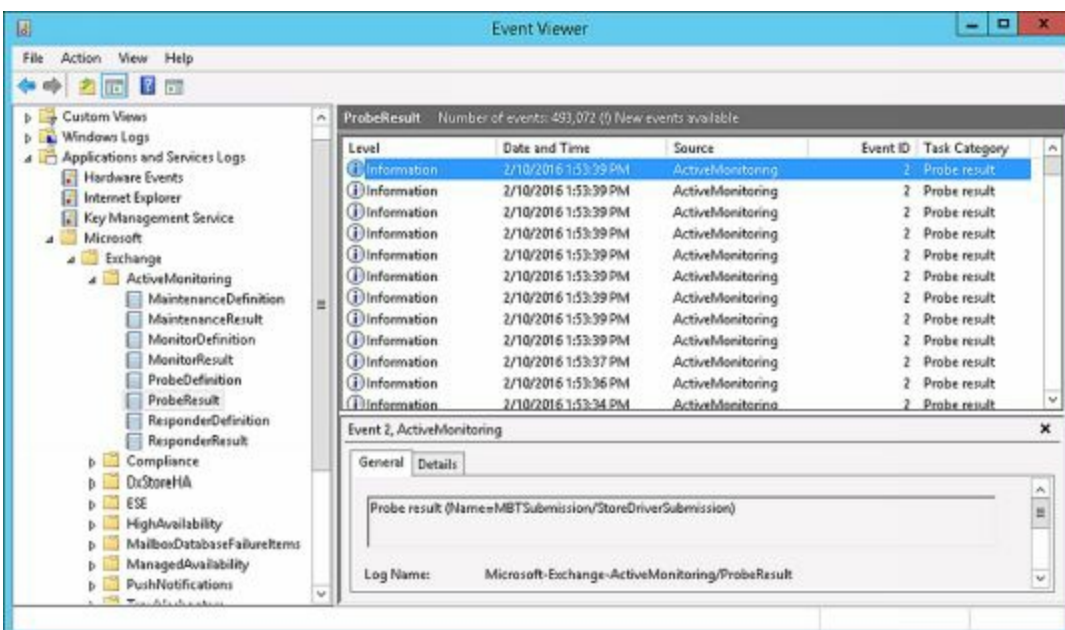
Managed availability is implemented using:

- **Exchange Health Manager Worker process (MSExchangeHMWorker.exe)** A working process that performs the runtime management tasks.
- **Exchange Health Manager Service (MSExchangeHMHost.exe)** A controller process used to execute and manage the work process. If the worker process becomes nonresponsive or otherwise fails, the controller process is used to recover the worker process.



During startup, the health manager worker process reads XML configuration files and initializes the probes, monitors, and responders used by the managed availability framework. The worker process stores runtime data in the registry and writes results to the event logs as well. In Event Viewer, you can find these logs by navigating to Applications and Services Logs > Microsoft > Exchange > ManagedAvailability.

As discussed further in Chapter 30, "Troubleshooting Exchange Server 2016," in the "Tracking Server Health" section, you can use Get-HealthReport and Get-ServerHealth to check the state and health of Exchange resources. Each tracked resource has customized sets of probes, monitors, and responders that help to ensure its availability. Probe definitions identify the Exchange resource to track and the time interval in which the resource is checked. Monitor definitions identify the specific state of the resource based on the collected data. In Event Viewer, you can find definitions and results for probes, monitors, and responders under Applications and Services Logs > Microsoft > Exchange > ActiveMonitoring.



Exchange tracks the transition state internally by using the TargetHealthState property

associated with a responder, where

- 0 indicates an alert threshold is no longer met
- 1 indicates a healthy state
- 2 indicates a degraded state
- 3 indicates an unhealthy state
- 4 indicates an unrecoverable state
- 5 indicates a Degraded1 state
- 6 indicates a Degraded2 state
- 7 indicates an Unhealthy1 state
- 8 indicates an Unhealthy2 state
- 9 indicates an Unrecoverable1 state
- 10 indicates an Unrecoverable2 state

When a resource transitions from one state to another is determined by the monitor definition. As soon as the monitor engine detects an unhealthy or degraded state for a responder, the transition state of that resource is shown as Unhealthy or Degraded respectively, which will trigger a recovery action. Whether a resource is shown as Unhealthy or Degraded depends on the data collected. For example, if a resource is unavailable, the resource may be listed as Unhealthy. If a resource is available but is slow to respond due to high latency or a high level of activity, the resource may be listed as Degraded.

Once in an Unhealthy state, the health manager may transition the resource to another state. For example, after 30 seconds in an unhealthy state, the resource may be transitioned to an Unhealthy1 state. After 300 seconds in an unhealthy state, the resource may be transitioned to an Unhealthy2 state. After 3000 seconds in an unhealthy state, the resource may be transitioned to an Unrecoverable state. Once in an Unrecoverable state, the health manager may transition the resource through the related Unrecoverable, Unrecoverable1, and Unrecoverable2 states.

Responder definitions detail the specific recovery actions to take based on the transition state of the Exchange resource. The exact response to an unhealthy state depends on the affected resource. Although the initial response to an unhealthy state might be to restart the related service or application pool, a subsequent response might be to restart the server, and a final response might be to take the server offline so that it no longer accepts traffic.

Creating and Managing Database Availability Groups

Database availability groups are a container in Active Directory and a logical layer on top of Windows Clustering. You can create and manage database availability groups in a variety of ways. Establishing a database availability group and making it operational requires the following at a minimum:

1. Preparing for deployment
2. Creating a database availability group
3. Adding member servers to the group
4. Designating a witness server
5. Creating an availability group network

These tasks and general management tasks for database availability groups are discussed in the sections that follow.

Preparing for DAGs

A database availability group defines a set of servers that provide automatic database-level recovery from database failures. Only members of the Organization Management group or the Database Availability Groups Role can create database availability groups. Only members of the Organization Management group or the Database Copies Role can manage mailbox database copies.

When you create a database availability group, you can specify a witness server or let Exchange choose one for you. The witness server's role is to help maintain the state of the group. It does this by maintaining the quorum when there is an even number of members in the group. On the witness server, you can designate a directory, called the witness directory, for use by the database availability group, or you can let Exchange create a default directory for you. By default, the witness directory is created as a subdirectory of %SystemDrive%\DAGFileShareWitnesses with the name set the same as the fully qualified domain name of the DAG.

Exchange creates and secures the witness directory automatically as part of configuring the witness server for use. The witness directory should not be used for any purpose other than for the database availability group witness server. The requirements for the witness server are as follows:

- The witness server cannot be a member of the database availability group.
- The witness server must be in the same forest as the database availability group.
- The witness server must be running a current version of Windows Server.

To be sure that Exchange administrators are aware of the availability of the witness server and that the server remains under the control of an Exchange administrator,

Microsoft recommends using an Exchange 2016 server to host the witness directory. Using an Exchange 2016 server as the witness also ensures that Exchange has sufficient permissions to remotely create and share the witness directory. The preferred witness server is a Mailbox server in the same Active Directory site as the majority of the members of the database availability group.

A single server can act as a witness for multiple database availability groups; however, every database availability group must have a separate witness directory. When your organization has multiple datacenters, placement of the witness server is especially important. While your organization may have two datacenters, a third location with a separate network infrastructure is needed for automatic failover from a site-level event. If your organization has a datacenter pair but doesn't have an alternate location with separate network infrastructure, you should:

- Consider placing the witness server in an Azure deployment, which would then be your alternate location for ensuring site failover is automatic.
- Or place the witness server in the datacenter where the majority of users are physically located and also make sure the Primary Active Manager for each availability group is located in this same datacenter.


The witness directory doesn't need to be fault tolerant and doesn't require any other special considerations. If you need to reset permissions on the witness directory or recreate the witness directory in its original location, you can use Set-DatabaseAvailabilityGroup as long as the cluster quorum is intact.

NOTE Cluster quorum ensures consistency of the DAG. Quorum represents a shared view of members and resources and also is used to describe the shared physical configuration within the DAG. Having quorum ensures that only one subset of cluster members is functioning in the DAG.

TIP Ideally, you'll locate the witness server in the same datacenter as DAG members. Although a server cannot act as a witness server for a DAG of which it is a member, a DAG member can act as a witness server for another DAG.

If the witness server isn't running Exchange 2016, Exchange 2013 or Exchange 2010, you must add the Exchange Trusted Subsystem group to the local Administrators group on the witness server prior to creating the DAG. Adding this group ensures that Exchange 2016 can create and share the witness directory. To add the Exchange Trusted Subsystem group to the local Administrators group, follow these steps:


1. In Control Panel, select **User Accounts**, and then select **Give Other Users Access To This Computer**.



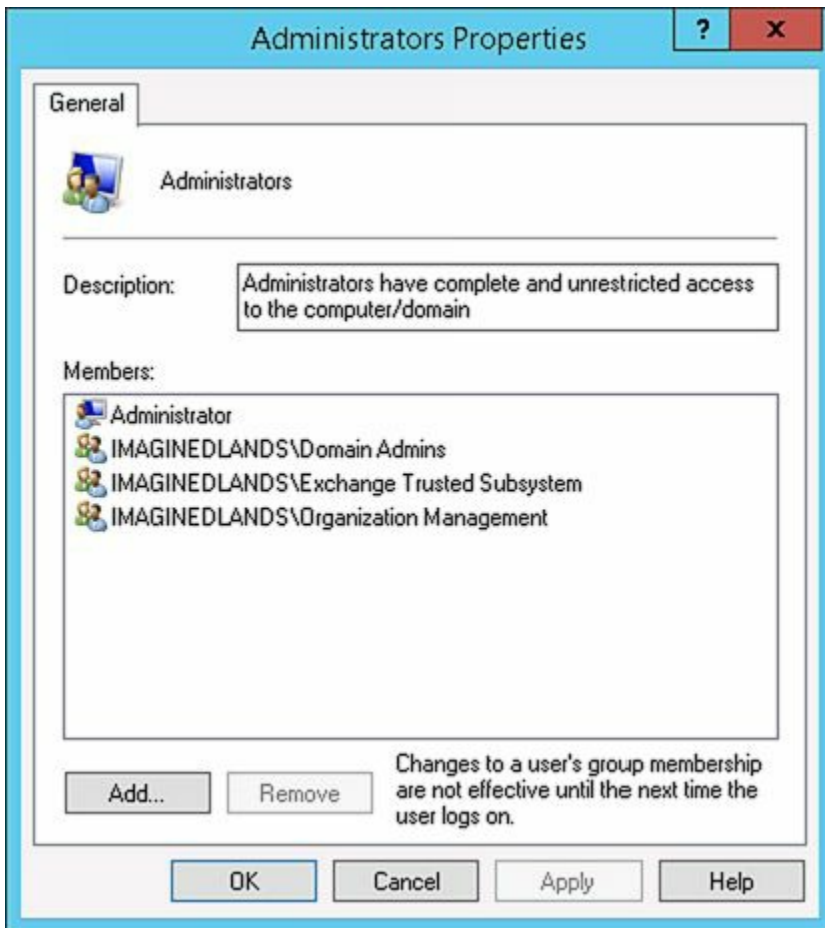
2. In the User Accounts dialog box, on the Advanced tab, click **Advanced**.



3. In the Local Users And Groups console, select **Groups**, and then double-click **Administrators**.



4. In the Administrators properties dialog box, select **Add**.
5. In the Select Users, Computers, Service Accounts, Or Groups dialog box, enter **Exchange Trusted Subsystem** and then click OK. The Exchange Trusted Subsystem is then added to the local Administrators group.



Beginning with Exchange 2016, you can now deploy a DAG with or without an administrative access point for the failover cluster. A failover cluster for the DAG is created without an administrative access point by default. Although a recommended best practice, this configuration is only supported when all member servers are running Windows Server 2012 R2 or later. Alternatively, you can specify one or more static IPv4 addresses for the DAG to use for dynamic IP addressing.

There are important differences between DAGs that have or don't have administrative access points for their failover clusters:

- A DAG with a cluster administrative access point uses a cluster name object (CNO). The CNO must be pre-staged and provisioned by assigning appropriate permissions.
- A DAG without a cluster administrative access point does not have a cluster name object (CNO). This also means you don't need to pre-stage or pre-provision a CNO.

To set up the database availability group, Exchange creates an msExchMDBAvailabilityGroup object and related objects in Active Directory Domain Services (AD DS). These objects represent the database availability group, its members, networks, and attributes. The msExchMDBAvailabilityGroup directory object is used to store information about the database availability group, such as server membership information. Information about the included databases is stored in the cluster database. When you add the first server to a database availability group, a failover cluster is automatically created for the database availability group and failover monitoring is initiated. The failover cluster heartbeat mechanism and cluster database are then used to track and manage information about the database availability group.

After a database availability group has been created, you can add or remove member servers. When the first Mailbox server is added to a database availability group, the following occurs:

- The Windows Failover Clustering component and related management tools are installed, if they are not already installed.

IMPORTANT Windows Failover Clustering is available on servers that are running Windows Server 2012, Windows Server 2012 R2 and Windows Server 2016. Each Mailbox server in the database availability group should have at least two network interface cards in order to have separate replication and messaging networks.

- A failover cluster is created using the name of the database availability group. For the purposes of authentication and access permissions, the cluster is represented by a computer account that is created in the default container for computers. This computer account is referred to as the cluster virtual network name account or the cluster network object.
- The server is added to the msExchMDBAvailabilityGroup object in Active Directory. When you create a database availability group that has an administrative access point, an IP address is assigned to the group. Then when you add the first server to the group, the name and IP address of the database availability group are registered in Domain Name System (DNS) using a Host (A) record. The name must be no longer than 15 characters and must be unique within the Active Directory forest. Otherwise, when you create a group without an administrative access point, only a corresponding Cluster Network Object (CNO) is created in Active Directory with the name of the DAG.
- The cluster database is updated with information about the databases that are mounted on the server.
- Exchange examines the current network configuration, as presented by the cluster. If the server has a properly configured network card, the configuration of that network card is used to create the replication network. If the server has two network cards, the configuration settings of those network cards are used to create separate replication and messaging networks.
- A base directory is created on the witness server. If you specified a directory during DAG creation, this directory is created. Otherwise, the %SystemDrive%\DAGFileShareWitnesses directory is created. Permissions are set so that the local Administrators group has full control.

NOTE The witness directory and witness file share aren't created until needed. Permissions are set so that the network name account representing the cluster has full control.

When you add the second and subsequent servers to the DAG, the following occurs:

- The server is joined to the failover cluster for the DAG.
- The server is added to the msExchMDBAvailabilityGroup object in Active Directory

Domain Services.

- The cluster database is updated with information about the databases that are mounted on the server.

When a database availability group has a single member server, the failover cluster initially uses the Node Majority quorum mode. When you add the second Mailbox server to the database availability group, Exchange changes the cluster quorum to the Node and File Share Majority quorum model and begins by using the Universal Naming Convention (UNC) path and directory for the cluster quorum. If the witness directory does not exist, Exchange automatically creates it at this point and configures its security with full control permissions for local administrators and the cluster network computer account for the database availability group.

REAL WORLD Every failover cluster has a resource that is responsible for maintaining the witness logs. This resource is called the quorum or witness resource. The quorum resource writes information about all cluster database changes to the witness logs, ensuring that the cluster configuration and state data can be recovered. When you create a database availability group, Exchange automatically determines the appropriate quorum configuration for your cluster based on the number of member servers. When a DAG has an odd number of members, Exchange uses the Node Majority quorum model. When a DAG has an even number of members, Exchange uses a Node and File Share Majority quorum model.

In a Node Majority cluster configuration, servers have a local quorum device. This device stores the cluster configuration information. In a Node and File Share Majority cluster configuration, servers use a witness file share rather than a quorum (witness) device. Otherwise, the Node and File Share Majority configuration works like the Node Majority configuration. The change from one model to the other should happen automatically. If it doesn't, run `Set-DatabaseAvailabilityGroup` with only the `-Identity` parameter, which will update the quorum settings for the DAG.

Before you create a database availability group that will use an administrative access point, you should pre-stage and prepare the cluster name object. You pre-stage the cluster name object by creating a computer account that will be used as the cluster's name resource. The name resource is a Kerberos-enabled object that acts as the cluster's identity and provides the cluster's security context.

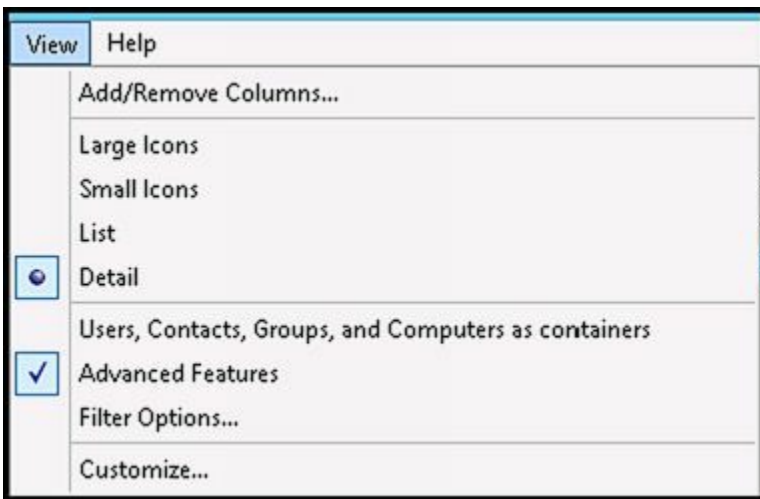


To pre-stage the cluster name object for a DAG with an administrative access point, complete the following steps:


1. In the Active Directory Users And Computers console tree, right-click the container in which you want to place the computer account, click New, and then click Computer.
2. In the Computer Name text box, enter the name that you want to use for the DAG. For example, if you are creating the first DAG in the Active Directory forest, you may want to enter **DAG01** as the name. The name can be up to 15 characters. The name must be unique in the Active Directory forest and cannot contain spaces or other special characters.
3. If Windows Deployment Services are not installed, click OK to create the computer account. Otherwise, click Next twice, and then click Finish.

Next, prepare the cluster name object for the DAG containing the administrative access point. You prepare the cluster name object by completing the following steps:

1. Open Active Directory Users And Computers. If Advanced Features aren't enabled, enable them by selecting **Advanced Features** on the View menu.



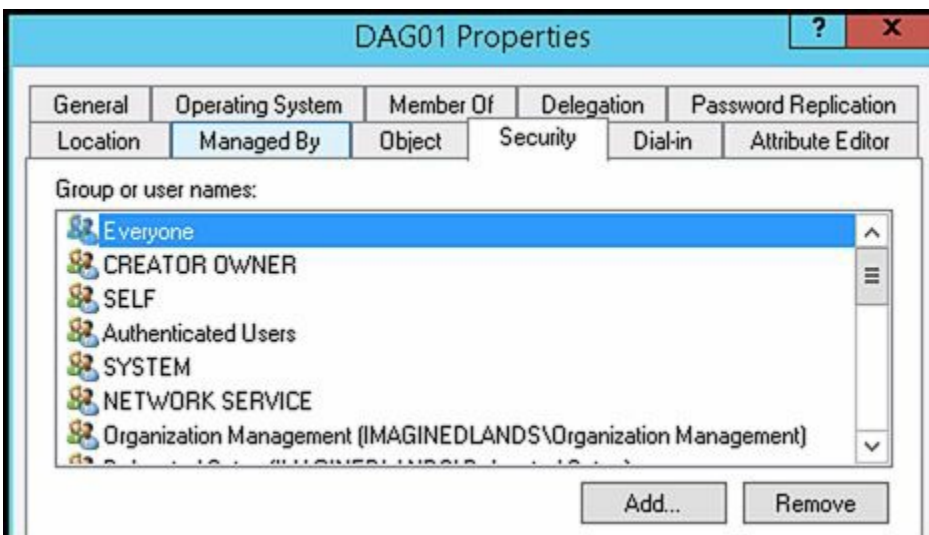
2. Right-click the computer account for the DAG, and then select **Disable Account**.



3. When prompted to confirm that you want to disable the account, select **Yes** and then select **OK**.

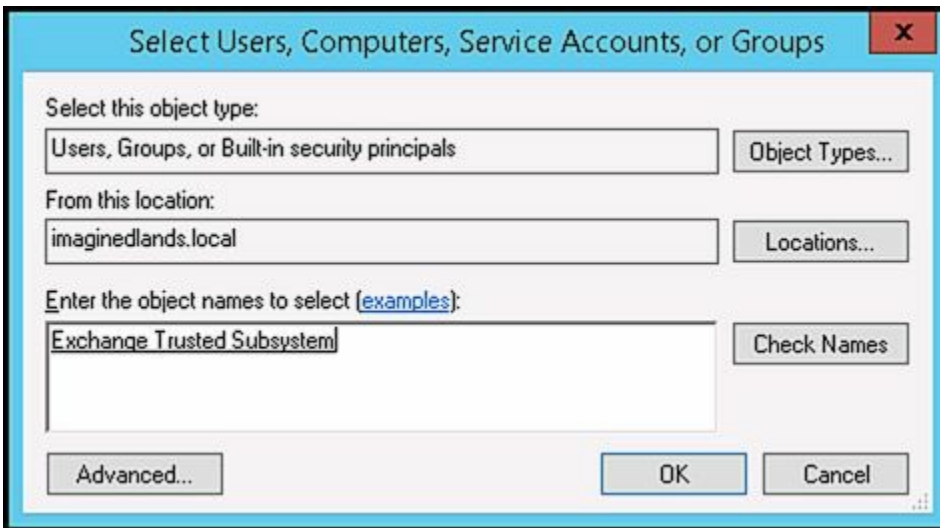


4. Right-click the computer account for the DAG, and then select **Properties**. In the properties dialog box, select the Security tab and then select **Add**.




5. In the Select User, Computer, Service Account, Or Group dialog box, enter **Exchange Trusted Subsystem** as the name of the group to which you want

to grant privileges, and then click OK.



6. With Exchange Trusted Subsystem selected in the Group Or User Names list, select **Full Control** in the **Allow** column, and then select **OK** to grant full control permissions on the cluster name object to the Exchange Trusted Subsystem.



SECURITY ALERT The Exchange Trusted Subsystem group has as members all the machine accounts for Exchange servers in the domain and as such, this group can be used to manage the cluster name object. Alternatively, you can enter the name of the first Mailbox server you are adding to the DAG and then grant full control permissions to the related computer object to ensure that the LOCAL SYSTEM security context on that server will be able to manage the cluster name object.

REAL WORLD When Windows Firewall is enabled, you must enable inbound exceptions for Windows Management Instrumentation (WMI) and File And Printer Sharing on the witness server. Keep in mind that if you don't specify a witness server, Exchange searches the local Active Directory site for an Exchange server that isn't a member of the DAG and configures this server as the witness server. To create the required inbound exceptions for Windows Firewall, follow these steps:

1. In Control Panel, select System And Security, and then select Windows Firewall.
2. In the left pane, select Allow An App Or Feature Through Windows Firewall.
3. If File And Printer Sharing is not selected for the Domain profile, select it under Allowed Apps And Features.
4. If Windows Management Instrumentation (WMI) is not selected for the Domain profile, select it under Allowed Apps And Features.

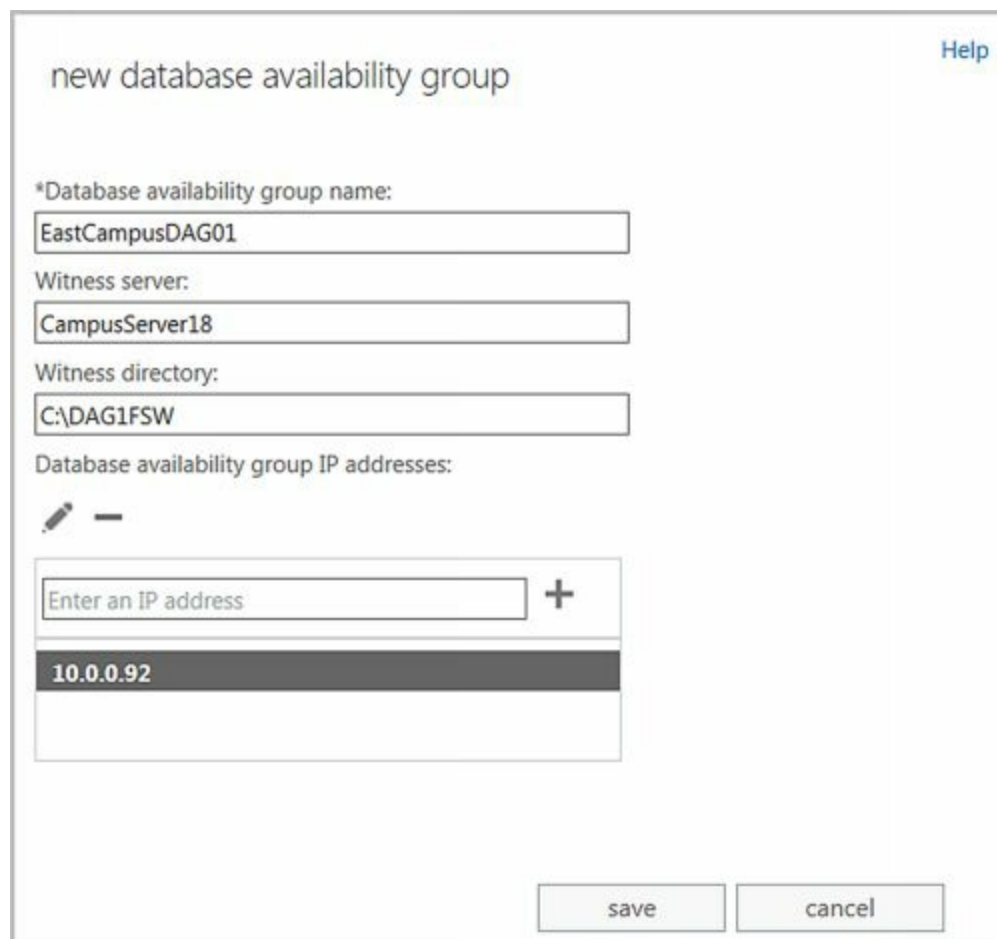
5. If you made changes to the Windows Firewall configuration, select OK.

If Windows Firewall is enabled and these exceptions are not created, you may see error messages warning that Exchange wasn't able to create the default witness directory or that Exchange is unable to access file shares on the witness server. You may see an error message stating: The network path was not found. Or you may see an error message stating: WMI exception occurred on the server. The RPC server is unavailable.

Creating Database Availability Groups

Once you've performed any necessary preparatory steps, you can create the database availability group by completing the following steps:

1. In the Exchange Admin Center, select Servers in the Navigation menu, and then select Database Availability Groups.
2. Select the New button to create the DAG. You should now see the New Database Availability Group dialog box, as shown in Figure 18-1.



The screenshot shows a dialog box titled "new database availability group" with a "Help" link in the top right corner. The dialog contains the following fields and controls:

- *Database availability group name:
- Witness server:
- Witness directory:
- Database availability group IP addresses: (with a "+" button to add more addresses)
- Buttons: "save" and "cancel"

FIGURE 18-1 Set the database availability group name and file locations.

3. In the Database Availability Group Name text box, enter the name of the pre-staged computer account for the DAG.
4. Optionally, provide the name of a server in the same Active Directory forest as the DAG to act as the witness server. Because this server cannot be a member of the database availability group, be sure that you don't select servers that will be

members of the database availability group you are configuring. The witness server must be running a current version of Windows Server, but doesn't need to be running Microsoft Exchange. However, if the witness server isn't running Microsoft Exchange, you'll need to modify the local Administrators group on the witness server prior to creating the DAG and add the Exchange Trusted Subsystem security group as a member.

NOTE The server you select as the witness server can be a member of a different database availability group. Also note that if you don't specify a witness server, Exchange attempts to automatically select a witness server by looking in the same Active Directory site as the majority of the DAG members for a Mailbox server that isn't part of the DAG.

5. Optionally, provide the local folder path for a directory that will be used to store witness data, such as C:\WitnessDir. If the directory does not exist, Exchange attempts to create it for you on the witness server. If you don't specify a witness directory, Exchange attempts to create a directory named relative to the database availability group on the witness server's system drive.

NOTE Exchange must have appropriate permissions on the server to create and then share the witness directory. Although you can set the local directory path, the share name is set automatically in the form DAGName.DomainName, such as WestCampusDag1.IMAGINEDLANDS.COM. This share is configured so that the cluster name object has full control.

TIP As long as the witness server is an Exchange server in the same forest, Exchange should be able to create and share the directory. If Exchange is unable to create and share the directory, you'll see an error message and will need to take appropriate corrective actions. You can use the Set-DatabaseAvailabilityGroup with the -WitnessDirectory parameter to specify a new directory to use at any time. You also can set a new directory by double-clicking the DAG in the Exchange Admin Center, entering a new directory path in the Witness Directory field, and then clicking OK.

If the witness server is not an Exchange 2016 server, you have to add the Exchange Trusted Subsystem security group to the local Administrators group on the witness server.

6. IP address assignment controls whether the DAG will have an administrative access point:
 - To create a DAG without an administrative access point, either provide no IPv4 address or enter 255.255.255.255 as the IPv4 address. Remember, all DAG members must be running Windows Server 2012 R2 or later.
 - To create a DAG with an administrative access point and static IP addressing, enter an IP address to use, and then select Add. Repeat this process to specify other static IP addresses to use.
 - To create a DAG with an administrative access point and dynamic IP addressing,

enter 0.0.0.0 as the IP address to use.

7. Select Save to create the database availability group. If an error occurred, you'll need to take the appropriate corrective action. Otherwise, you can now add member servers to the database availability group.

In Exchange Management Shell, you can create database availability groups by using the `New-DatabaseAvailabilityGroup` cmdlet. Listing 18-1 provides the syntax and usage. Set DAG name using the `-Name` parameter. To control whether the DAG has an administrative access point and IPv4 addressing, do the following:

- Create a DAG without an administrative access point by using the value `255.255.255.255` for the `-DatabaseAvailabilityGroupIpAddresses` parameter. Ensure all DAG members are running Windows Server 2012 R2 or later.
- Create a DAG with an administrative access point and static IP addressing by setting the `-DatabaseAvailabilityGroupIpAddresses` parameter to the specific IPv4 addresses to use.
- Create a DAG with an administrative access point and dynamic IP addressing by setting the `-DatabaseAvailabilityGroupIpAddresses` parameter to `0.0.0.0`.

NOTE Don't confuse the local witness directory with the witness file share. The local witness directory has a local file path on the witness server, such as `C:\WitnessShare`. When you specify the witness directory, Exchange creates the base directory and then creates and shares a subdirectory within this directory as appropriate.

LISTING 18-1 `New-DatabaseAvailabilityGroup` cmdlet syntax and usage

Syntax

```
New-DatabaseAvailabilityGroup -Name DAGName  
[-DatabaseAvailabilityGroupIpAddresses IPAddress1, IPAddress2, IPAddressN ]  
[-WitnessServer ServerName]  
[-WitnessDirectory LocalDirOnWitnessServer]  
[-DomainController FullyQualifiedName]  
[-ThirdPartyReplication <Disabled | Enabled>]
```

Usage

```
New-DatabaseAvailabilityGroup -Name "EastCampusDAG1"  
-WitnessServer "WinServer19"  
-WitnessDirectory "C:\EastCampusDAG1"
```

```
New-DatabaseAvailabilityGroup -Name "WestCampusDAG1"  
-WitnessServer "WinServer19"  
-WitnessDirectory "C:\WestCampusDAG1"  
-DatabaseAvailabilityGroupIpAddresses 255.255.255.255
```

```
New-DatabaseAvailabilityGroup -Name "NorthCampusDAG1"  
-DatabaseAvailabilityGroupIpAddresses 0.0.0.0
```

Managing Availability Group Membership

When you add a server to a database availability group, the server works with the other servers in the group to provide automatic, database-level recovery from database, server, and network failures. When member servers have only one network adapter card, the DAG uses the same network for both messaging and replication traffic. When member servers have two network cards, the DAG uses one network card primarily for messaging traffic and the other network card is typically dedicated to replication traffic. If you add more than two network cards to member servers, these additional network cards can be configured for replication, giving the DAG additional replication networks to handle increased workloads.

NOTE Member servers in a DAG can have zero or more replication networks but only one messaging network. For optimal operation, servers should have at least two network interface cards with each network interface card configured to use a different subnet.

Keep the following in mind when planning database availability group membership:

- All servers in a DAG must be running the same operating system, which can be Windows Server 2012, Windows Server 2012 R2 or Windows Server 2016.
- When you add the first Mailbox server to a database availability group, the group must be assigned an IP address. If no IP address is assigned, Exchange uses DHCP to obtain an IP address for the group.
- If you no longer want a server to be a member of a group, you can remove it from the group and the server will no longer be automatically protected from failures. Keep in mind that you must remove all replicated database copies from the server before you can remove it from the group.

You can add a Mailbox server to or remove a Mailbox server from a database availability group by completing the following steps:

1. In the Exchange Admin Center, select Servers in the Navigation menu, and then select Database Availability Groups to view existing availability groups, as shown in Figure 18-2.

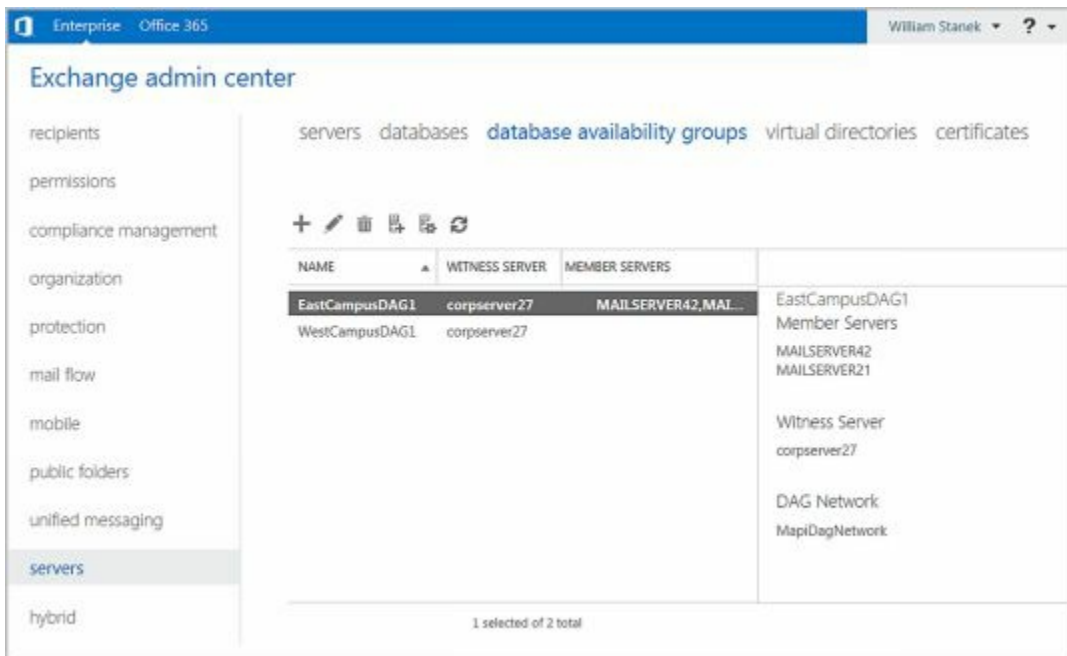


FIGURE 18-2 View configured database availability groups.

2. Select the DAG you want to configure, and then select the Manage DAG Membership button. In the Manage Database Availability Group Membership dialog box, shown in Figure 18-3, any current DAG members are listed by name. You can now:
 - Click the Add button to add a server to the database availability group. In the Select Server dialog box, select a server, click Add, and then repeat as necessary to select other servers.
 - Select a server and then click the Remove button to remove a server from the database availability group.
3. When you are finished selecting servers, choose OK and then choose Save. For each server you added, Exchange Admin Center will install the required Windows Failover Clustering components, and then add the server to the DAG. Subsequently, Exchange Admin Center will create and configure the witness directory and file share. For each server you removed, Exchange Admin Center will attempt to remove the server from the DAG. If an error occurs during these tasks, you will need to take the appropriate corrective action; otherwise, click Close when these tasks have completed successfully.

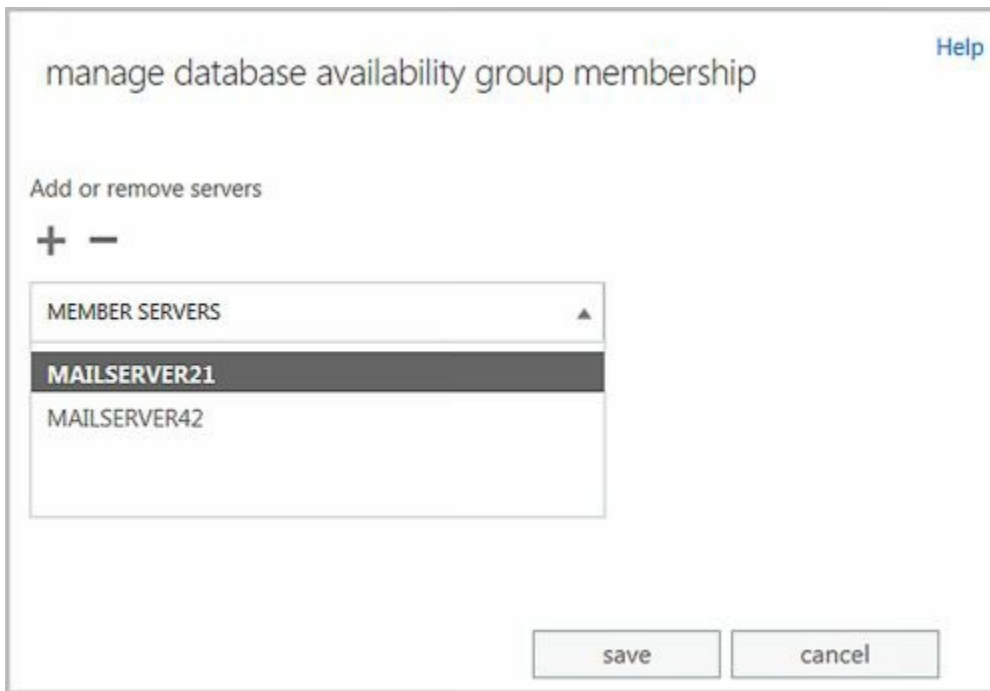


FIGURE 18-3 Add group members.

In Exchange Management Shell, you can list database availability groups by using `Get-DatabaseAvailabilityGroup`. If you enter `Get-DatabaseAvailabilityGroup` without additional parameters, you'll see a list of all availability groups in the current Active Directory forest as well as the member servers and operational servers for those groups, as shown in the following example and sample output:

```
Get-DatabaseAvailabilityGroup
```

Name	Member Servers	Operational Servers
EastCampusDAG1	MailServer42, MailServer21	MailServer42, MailServer21
WestCampusDAG1	MailServer44, MailServer96	MailServer44, MailServer96

Use the `-Identity` parameter to specify the name of the database availability group to query. Add `-Status` to any query to include real-time status information.

You add or remove group members by using `Add-DatabaseAvailabilityGroupServer` and `Remove-DatabaseAvailabilityGroupServer`. Listings 18-2 and 18-3 provide the syntax and usage.

LISTING 18-2 Add-DatabaseAvailabilityGroupServer cmdlet syntax and usage

Syntax

```
Add-DatabaseAvailabilityGroupServer -Identity DAGName
-MailboxServer ServerToAdd [-DomainController FullyQualifiedName]
[-SkipDagValidation {$true | $false}]
```

Usage

```
Add-DatabaseAvailabilityGroupServer -Identity "EastCampusDAG1"
-MailboxServer "MailServer62"
```

LISTING 18-3 Remove-DatabaseAvailabilityGroupServer cmdlet syntax and usage

Syntax

```
Remove-DatabaseAvailabilityGroupServer -Identity DAGName  
-MailboxServer ServerToRemove [-ConfigurationOnly <$true | $false>]  
[-DomainController FullyQualifiedName] [-SkipDagValidation {$true|$false}]
```

Usage

```
Remove-DatabaseAvailabilityGroupServer -Identity "EastCampusDAG1"  
-MailboxServer "MailServer62"
```

Managing Database Availability Group Networks

Each database availability group should have a minimum of two networks: one for replication traffic, referred to as the group's *replication network*, and one for MAPI and other traffic, referred to as the group's *messaging network*. Although a DAG can have only one messaging network, you can create additional replication networks in a database availability group and configure them by using the Exchange Management tools. Having multiple replication networks helps scale the DAG to meet increasing requirements.

REAL WORLD Although highly available servers have multiple logical networks, you can use a single, non-teamed network interface to handle both replication network and messaging network traffic. This configuration is in fact recommended as part of the preferred architecture to simplify the network stack and remove the requirement to manually eliminate heartbeat cross-talk.

By default, Exchange 2016 automatically creates DAG networks based on the configuration of network adapter cards installed on member servers. If a DAG member has multiple network cards and those cards are configured on separate networks, Exchange normally will configure the DAG members with one messaging network and one or more dedicated replication networks automatically. You can manually configure DAG networks as well but must first disable automatic network configuration.

You can enable manual network configuration for a DAG by completing the following steps:

1. In the Exchange Admin Center, select Servers in the Navigation menu, and then select Database Availability Groups to view existing availability groups.
2. Double-click the DAG you want to manually configure.
3. In the properties dialog box, on the General page, select the Configure Database Availability Group Networks manually check box, and then select Save.
4. You can now manually configure and manage the networks for the DAG. Keep in mind that if you later disable manual configuration, any manually created networks and related settings will be removed and Exchange Admin Center will create new DAG networks based on the current configuration of DAG members.

NOTE Use the `-ManualDagNetworkConfiguration` parameter of the `Set-DatabaseAvailabilityGroup` cmdlet to enable a manual network configuration. Set

the parameter to \$true to enable or \$false to disable manual network configuration.

Once you enable manual network configuration, you can manually create and manage network settings for the DAG. Each database availability group network must have a unique name of up to 128 characters, one or more subnet associations, and an optional description of up to 256 characters. When you configure a network, you can dedicate the network to replication traffic or dedicate the network to MAPI traffic.

NOTE Disabling replication does not guarantee that Exchange will not use a network for replication. If all configured replication networks are offline, failed, or otherwise unavailable, and only a nonreplication network remains, Exchange will use that network for replication until a replication-enabled network becomes available.

REAL WORLD Every network address has a network identifier that identifies the network and a host identifier that identifies the individual host on the network. The network ID is seen as the prefix of an IPv4 or IPv6 address, and the host ID is the suffix. When you define an availability group network, you need to identify the network and then specify the number of bits in the network number that are part of the network ID (and the remaining bits are understood to be part of the host ID). To write a block of IPv4 addresses and specify which bits are used for the network ID, you write the network number followed by a forward slash and the number of bits in the network ID, as follows:

NetworkNumber/# of bits in the network ID

The slash and the number of bits in the network ID are referred to as the network prefix. By default, Class A IPv4 networks have 8 bits in the network ID, Class B IPv4 networks have 16 bits, and Class C IPv4 networks have 24 bits.

IPv6 doesn't use subnet masks to identify which bits belong to the network ID and which bits belong to the host ID. Instead, each IPv6 address is assigned a subnet prefix length that specifies how the bits in the network ID are used. The subnet prefix length is represented in decimal form. If 48 bits in the network ID are used, the subnet prefix length is written as FEC0:1234:5678::/48 to represent the IPv6 addresses FEC0:1234:5678:: through FEC0:1234:5678::FFFF:FFFF:FFFF:FFFF.

You can create a network for a database availability group by completing the following steps:

1. In the Exchange Admin Center, select Servers in the Navigation menu, and then select Database Availability Groups to view existing availability groups.
2. Select the DAG you want to configure and then select the New DAG Network button.
3. In the New Database Availability Group Network dialog box, shown in Figure 18-4, enter a unique name for the database availability group network of up to 128 characters and then provide an optional description for the database availability group network of up to 256 characters.

FIGURE 18-4 Create a network for the availability group.

4. Under Subnets, click Add to add a network subnet to the database availability group network. Subnets should be entered by using a format of *IPv4Address/Bitmask*, such as 192.168.15.0/24, or *IPv6Address/NetworkSubnetPrefix*, such as FEC0:1234:5678::/48. The subnet must match the subnet used by one or more DAG members. If you add a subnet that is currently associated with another database availability group network, the subnet is removed from the other database availability group network and associated with the network being created.
5. Click Save. If an error occurred, you need to take the appropriate corrective action before you can create the network. If a warning is displayed, Exchange Admin Center will create the network but the network might not be operational until you correct the problem that prompted the warning. Otherwise, click Close when the task completes.

When the DAG is selected in Exchange Admin Center, the details pane lists the networks associated with the DAG. If manual configuration of networks is enabled, you'll see options for managing each network in the details pane, as shown in Figure 18-5, and these options include:

- **Disable Replication** Configures the network with a preference for messaging; however, the DAG will use the network for replication if necessary. Note also that a DAG can have only one dedicated messaging network.
- **Remove** Removes a DAG network, providing the network doesn't have any active subnets. Before you can remove a network with active subnets, you must assign the subnets to other networks.
- **View Details** Opens the properties dialog box for the network. You can use the options in this dialog box to change the network name, network description, and

associated subnets. By selecting or clearing the Enable Replication checkbox, you can enable or disable replication on the network.

The properties dialog box for a DAG network also shows the status of subnets and network interfaces. Subnets and interfaces listed as Up are active. Subnets and interfaces listed as Down are inactive.

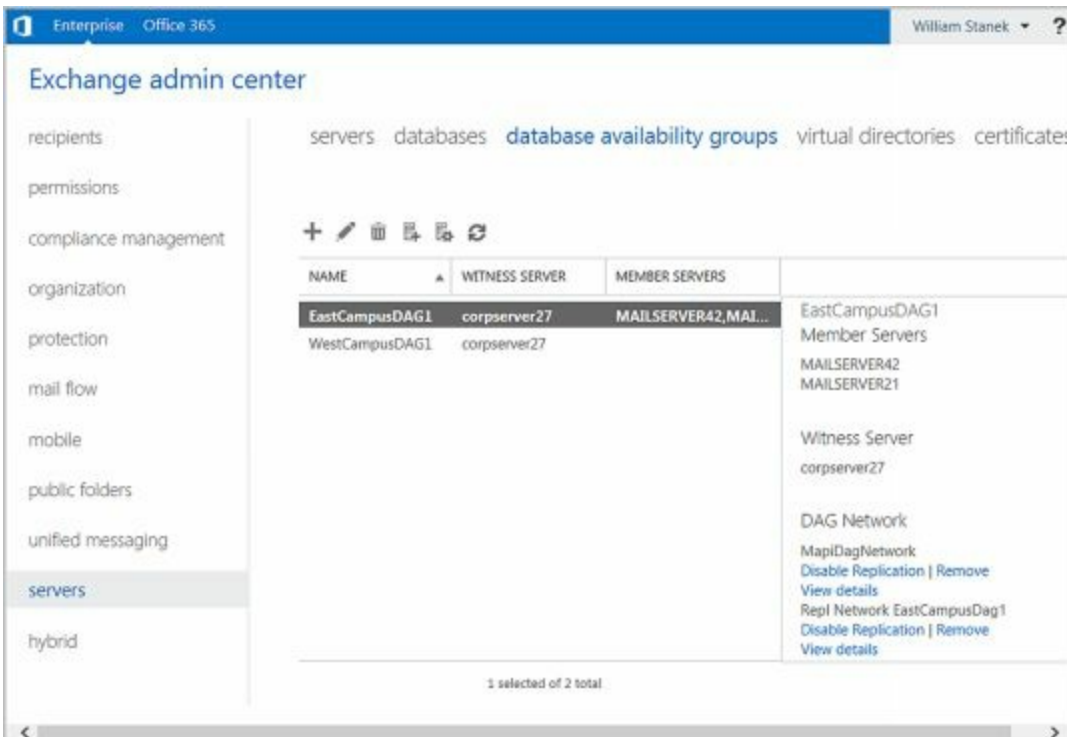


FIGURE 18-5 View the networks configured for a DAG.

In Exchange Management Shell, you can list availability DAG networks and their status by using `Get-DatabaseAvailabilityGroupNetwork`. If you enter `Get-DatabaseAvailabilityGroupNetwork` without additional parameters, you see a list of all configured networks for all availability groups. Use the `-Identity` parameter to specify the name of the network to query. Use the `-Server` parameter to obtain health information for the network from a specific Mailbox server. The following example lists detailed information for all the networks associated with `EastCampusDAG1`:

```
Get-DatabaseAvailabilityGroupNetwork -Identity EastCampusDAG1 | fl
```

The following example lists detailed information for network associated with `EastCampusDAG1` that have names starting with `Repl`:

```
Get-DatabaseAvailabilityGroupNetwork -Identity EastCampusDAG1\Repl* | fl
```

The detailed information is helpful as it lists the status of associated subnets and interfaces as shown in the following sample:

```
Name           : Repl Network EastCampusDAG1
Description    :
Subnets       : {{{10.0.0.0/24,Up}}}
Interfaces     : {{{MailServer21,Up,10.0.0.50}},
                {MAILSERVER42,Up,10.0.0.60}}
MapiAccessEnabled : True
```


ReplicationEnabled : True
IgnoreNetwork : False
Identity : EastCampusDAG1\Repl Network EastCampusDAG1
IsValid : True

You create or remove group networks by using `New-DatabaseAvailabilityGroupNetwork` and `Remove-DatabaseAvailabilityGroupNetwork`. Listings 18-4 and 18-5 provide the syntax and usage.

LISTING 18-4 `New-DatabaseAvailabilityGroupNetwork` cmdlet syntax and usage

Syntax

```
New-DatabaseAvailabilityGroupNetwork -Name NetworkName  
-DatabaseAvailabilityGroup DAGName  
[-Description Description] [-DomainController FullyQualifiedName]  
[-IgnoreNetwork <$true | $false>] [-ReplicationEnabled <$true | $false>]  
[-Subnets SubnetIds]
```

Usage

```
New-DatabaseAvailabilityGroupNetwork -DatabaseAvailabilityGroup  
"EastCampusDAG1" -Name "Primary DAG Network" -Description ""  
-Subnets "{192.168.10.0/24, 192.168.15.0/24}" -ReplicationEnabled $true
```

LISTING 18-5 `Remove-DatabaseAvailabilityGroupNetwork` cmdlet syntax and usage

Syntax

```
Remove-DatabaseAvailabilityGroupNetwork -Identity NetworkName  
[-DomainController FullyQualifiedName]
```

Usage

```
Remove-DatabaseAvailabilityGroupNetwork  
-Identity "EastCampusDAG1\Primary DAG Network"
```

Changing Availability Group Network Settings

Database availability group networks have several properties that you can configure, including the network name, description, associated subnets, and replication status. The replication status determines whether the network is used as the replication network for the group or the messaging network for the group. When replication is enabled, the network is used as the replication network for the group. When replication is disabled, the network is used as the messaging network for the group.

When manual network configuration is enabled, you can manage the settings for a group network by completing the following steps:

1. In the Exchange Admin Center, select **Servers** in the Navigation menu, and then select **Database Availability Groups** to view existing availability groups.
2. When you select the DAG you want to work with, the details pane lists the

associated networks. Each network has a related set of management options. Select the **View Details** option for the network you want to configure.

DAG Network

MapiDagNetwork
[Disable Replication](#) | [Remove](#)
[View details](#)

Repl Network EastCampusDag1
[Disable Replication](#) | [Remove](#)
[View details](#)

3. In the properties dialog box for the network, enter a new name if desired and optionally change the network description.

Repl Network EastCampusDAG1

*Database availability group network name:

Description:

4. Each network must contain at least one subnet. Subnets must be added by using a format of *IPAddress/Bitmask*, such as 192.168.15.0/24, or *IPv6Address/NetworkSubnetPrefix*, such as FEC0:1234:5678::/48. Use the options provided to add, edit, or remove subnets for the network.

Subnets:

+ ✎ -

SUBNET	STATUS
10.0.0.0/24	Up

Network interfaces:

NETWORK INTERFACE	STATUS
10.0.0.50	Up
10.0.0.60	Up

5. To establish the network as the replication network for the group, select the **Enable Replication** check box. Otherwise, clear the check box to use the network as the messaging network for the group.
6. Click **Save** to apply your settings.



You can use `Set-DatabaseAvailabilityGroupNetwork` to configure basic settings for availability group networks. Listing 18-6 provides the syntax and usage for `Set-DatabaseAvailabilityGroupNetwork`.

LISTING 18-6 `Set-DatabaseAvailabilityGroupNetwork` cmdlet syntax and usage

Syntax

```
Set-DatabaseAvailabilityGroupNetwork -Identity NetworkName  
[-Description Description ] [-DomainController FullyQualifiedName]  
[-IgnoreNetwork <$true | $false>] [-Name NewName]  
[-ReplicationEnabled <$true | $false>] [-Subnets Subnets ]
```

Usage

```
Set-DatabaseAvailabilityGroupNetwork  
-Identity "EastCampusDAG1\Primary DAG Network"  
-ReplicationEnabled $False
```

Advanced options for the networks associated with availability groups are set at the group level. Advanced options you can configure include encryption, compression, and the TCP port used for replication. Database availability groups support data encryption by using the built-in encryption capabilities of the Windows Server operating system. When you enable encryption, database availability groups use Kerberos authentication between Exchange servers to encrypt and decrypt messages. Encryption helps maintain the integrity of the data. Network encryption is a property of the database availability group and not a property of a database availability group network.

You can configure database availability group network encryption by using the `-NetworkEncryption` parameter of the `Set-DatabaseAvailabilityGroup` cmdlet in Exchange Management Shell. The possible encryption settings are as follows:

- **Disabled** Network encryption is not used for any database availability group networks.
- **Enabled** Network encryption is used on all database availability group networks for replication and seeding.
- **InterSubnetOnly** Network encryption is used only with database availability group networks on the same subnet.
- **SeedOnly** Network encryption is used on all database availability group networks for seeding only.

Database availability groups also support built-in compression. You configure network compression by using the `-NetworkCompression` parameter of the `Set-DatabaseAvailabilityGroup` cmdlet in Exchange Management Shell. The possible compression settings are as follows:

- **Disabled** Network compression is not used for any database availability group networks.
- **Enabled** Network compression is used on all database availability group networks for replication and seeding.
- **InterSubnetOnly** Network compression is used only with database availability group networks on the same subnet.
- **SeedOnly** Network compression is used on all database availability group networks for seeding only.

You can specify the TCP port to use for replication by using the `–ReplicationPort` parameter of the `Set-DatabaseAvailabilityGroup` cmdlet in Exchange Management Shell.

Configuring Database Availability Group Properties

You can use the Exchange Admin Center or Exchange Management Shell to configure the properties of a database availability group, including the witness server and witness directory used by the database availability group. By using Exchange Management Shell, you can configure additional properties, such as encryption and compression settings, network discovery, the TCP port used for replication, alternate file share witness settings, and data center activation coordination mode.

To view or modify the properties of an availability group, complete the following steps:

1. In the Exchange Admin Center, select **Servers** in the Navigation menu, and then select **Database Availability Groups** to view existing availability groups.
2. Double-click the database availability group with which you want to work. This opens the Properties dialog box for the DAG.
3. On the **General** page, you'll see a list of member servers, the witness server's fully qualified domain name, and the location of the witness directory on the witness server.

EastCampusDAG1

general

IP address

Witness server:
corpserver27.pocket-consultant.com

Witness directory:
C:\DAG1FSW

Database availability group members:

NAME	IS OPERATIONAL
MAILSERVER21	Yes
MAILSERVER42	Yes

Configure database availability group networks manually

4. Using the Witness Server text box, you can specify a new witness server by

entering the fully qualified domain name of the new witness server. This server should be in the same Active Directory forest as the member servers and cannot be a current or future member of the database availability group.

5. Using the Witness Directory text box, you can specify a new witness directory on the witness server. If the directory does not exist, it will be created on the witness server.
6. Click **Save**.

In Exchange Management Shell, you can configure properties of database availability groups by using the `Set-DatabaseAvailabilityGroup` cmdlet. Listing 18-7 provides the syntax and usage.

LISTING 18-7 Set-DatabaseAvailabilityGroup cmdlet syntax and usage

Syntax

```
Set-DatabaseAvailabilityGroup -Identity DAGName  
[-DatabaseAvailabilityGroupIpAddresses IPAddresses]  
[-DatacenterActivationMode {"Off"|"DagOnly"}]  
[-DiscoverNetworks] [-DomainController FullyQualifiedName]  
[-NetworkCompression {"Disabled"|"Enabled"|"InterSubnetOnly"|"SeedOnly"}]  
[-NetworkEncryption {"Disabled"|"Enabled"|"InterSubnetOnly"|"SeedOnly"}]  
[-ReplicationPort TCPPort] [-AlternateWitnessServer ServerName]  
[-AlternateWitnessServerDirectory DirectoryPath]  
[-WitnessServer ServerName] [-WitnessServerDirectory DirectoryPath]
```

Usage

```
Set-DatabaseAvailabilityGroup -Identity "EastCampusDAG1"  
-NetworkCompression "Enabled" -NetworkEncryption "Enabled"  
-ReplicationPort 33898 -DatacenterActivationMode "Off"
```

Options for working with encryption, compression, and replication ports were discussed previously in “Changing Availability Group Network Settings.” Options that weren’t discussed include the datacenter activation coordinator mode, the alternate witness server, and alternate witness server directory. These options can be used as part of a datacenter switchover process. The alternate witness server must not be a part of the database availability group.

The data-center coordinator mode should be set for all database availability groups with three or more members that are extended to two or more physical locations. This mode cannot be enabled for groups with less than three members. When the datacenter coordinator is enabled, you can start, stop, and restore member servers in an availability group individually or collectively by using the following:

- **Start-DatabaseAvailabilityGroup** Activates member Mailbox servers in a recovered data center after a data-center switchover, as part of the failback process to the recovered data center. This command sets the configuration and state so that the

servers are incorporated into the operating database availability group and joined to the group's cluster. You use the `-MailboxServer` parameter to start a specific member server or the `-ActiveDirectorySite` parameter to start all members in a particular site.

```
Start-DatabaseAvailabilityGroup -Identity DAGName  
[-MailboxServer ServerName | -ActiveDirectorySite SiteName]  
[-ConfigurationOnly <$true | $false>]  
[-DomainController FullyQualifiedName]
```

NOTE You can also reactivate servers from a previously failed datacenter that has been restored to service. Before you can reactivate member Mailbox servers in a primary data center, the servers must first be integrated back into the operational database availability group. You reintegrate servers by running the `Start-DatabaseAvailabilityGroup` cmdlet and then using the `Move-ActiveMailboxDatabase` cmdlet to activate databases in the primary data center.

- **Stop-DatabaseAvailabilityGroup** Deactivates member Mailbox servers after a datacenter switchover. You use the `-MailboxServer` parameter to deactivate a specific member server or the `-ActiveDirectorySite` parameter to deactivate all members in a particular site.

```
Stop-DatabaseAvailabilityGroup -Identity DAGName  
[-MailboxServer ServerName | -ActiveDirectorySite SiteName]  
[-ConfigurationOnly <$true | $false>]  
[-DomainController FullyQualifiedName]
```

- **Restore-DatabaseAvailabilityGroup** Activates member Mailbox servers in a standby data center. Typically, this process is performed after the failure or deactivation of the active member servers in a primary data center. To activate all members in a particular site, you can use the `-ActiveDirectorySite` parameter.

```
Restore-DatabaseAvailabilityGroup -Identity DAGName  
[-ActiveDirectorySite SiteName]  
[-AlternateWitnessServer ServerName]  
[-AlternateWitnessDirectory DirectoryPath]  
[-DomainController FullyQualifiedName]  
[-UsePrimaryWitnessServer <$true | $false>]
```

Removing Servers from a Database Availability Group

Before you can remove a server from a database availability group, you must also remove all database copies from the server. To remove member servers from a DAG, select the DAG you want to manage, and then select the Manage DAG Membership button. In the Manage Database Availability Group Membership dialog box, select a server on the list of current members, and then select the Remove button. Repeat as necessary to remove members. Click Save to apply the changes. If an error occurs during these tasks, you will need to take the appropriate corrective action. Otherwise, click Close when these tasks have completed successfully.

After you remove the member servers, you can remove the database availability group

by selecting it and selecting the Delete button. When prompted to confirm, click Yes.

Removing Database Availability Groups

You can remove a database availability group only if it has no member servers. Therefore, before you can remove a database availability group, you must first remove any member servers from the group.

You can remove an empty availability group by completing the following steps:

1. In the Exchange Admin Center, select Servers in the Navigation menu, and then select Database Availability Groups to view existing availability groups.
2. On the Database Availability Group tab, select the database availability group you want to remove, and then select the Delete button.
3. When prompted to confirm the action, click Yes.

In Exchange Management Shell, you can remove database availability groups by using the Remove-DatabaseAvailabilityGroup cmdlet. Listing 18-8 provides the syntax and usage.

LISTING 18-8 Remove-DatabaseAvailabilityGroup cmdlet syntax and usage

Syntax

```
Remove-DatabaseAvailabilityGroup -Identity DAGName  
[-DomainController FullyQualifiedName]
```

Usage

```
Remove-DatabaseAvailabilityGroup -Identity "EastCampusDAG1"
```

Maintaining Database Availability Groups

The Microsoft Exchange Information Store service manages the active and passive databases configured on a Mailbox server. To improve performance, the service running on each server maintains a database cache of changes to active databases that haven't been applied to passive copies. In the event of a failover or switchover, the service can apply the changes in the cache to a passive copy and then make the passive copy the active copy. Most of the time, failover completes in about 30 seconds.

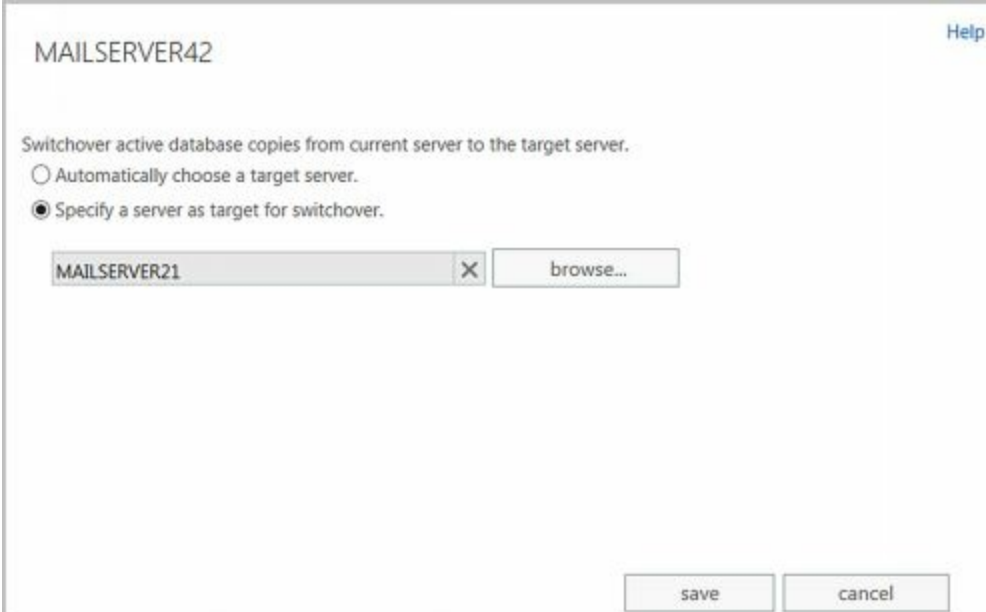
The difference between failover and switchover is important. When Exchange detects a failure of an active database, regardless of whether it is from database failure, server failure, or network failure, Exchange uses failover processes to mark the active database as inactive and dismount it and then mount and mark a passive database copy as the active copy. Prior to performing maintenance on a server or for testing or troubleshooting, you might want Exchange to switch from one database to another by marking an active database as inactive and then marking a passive database copy as the active copy.

Switching Over Servers and Databases

Failover and switchover occur at the database level for individual databases and at the server level for all active databases hosted by a server. When either a switchover or failover occurs, other Exchange 2016 server roles become aware of the switchover almost immediately and redirect client and messaging traffic automatically as appropriate.

You can switch over all active databases on a server by completing the following steps:

1. In the Exchange Admin Center, select Servers in the Navigation menu, and then select Servers. In the main pane, select the server that you are performing maintenance on, testing, or troubleshooting.



The screenshot shows a dialog box titled "MAILSERVER42" with a "Help" link in the top right corner. The main text reads: "Switchover active database copies from current server to the target server." Below this text are two radio button options: "Automatically choose a target server." (which is unselected) and "Specify a server as target for switchover." (which is selected). Under the selected option, there is a text input field containing "MAILSERVER21" and a "browse..." button to its right. At the bottom of the dialog box are two buttons: "save" and "cancel".

FIGURE 18-6 Switch over the active databases.

2. In the details pane, select Server Switchover. In the Server Switchover dialog box, shown in Figure 18-6, the default option is to allow Exchange to handle the switchover and select a server to take over the databases from the source server automatically. To accept the default, select Save. Otherwise, select Specify A Server As Target For Switchover and then select Browse. In the Select Server dialog box, select the server to take over, select OK. Keep in mind that you can select only a server that is already a member of the database availability group. You can't have copies outside the group either.
3. Select Save to apply the changes. When prompted to confirm the action, click Yes.

You can perform a switchover of an individual database by completing the following steps:

1. In the Exchange Admin Center, select Servers in the Navigation menu, and then select Databases.
2. In the main pane, select the database you want to work with. In the details pane, you see the available database copies, which are listed according to their copy status and health. Only the active copy will have a status of Active Mounted. All other database copies will display the current status of replication for the database copy, such as Passive Healthy.
3. Using the management options for the passive copy you want to activate, you can select View Details to display detailed information about the database copy, including the content index status, the overall status, the copy queue length, the replay queue length, and error messages.
4. To activate the copy, click the related Activate option. When you click Yes to confirm that you want to activate this copy, Exchange will dismount the current active mailbox database and establish the selected database copy as the active mailbox database.
5. If an error occurred, you need to take the appropriate corrective action before you can create the network. If a warning is displayed, Exchange Admin Center will create the network but the network may not be operational until you correct the problem that prompted the warning. Otherwise, click Close when the task completes.

When you are working with Exchange Management Shell, you can initiate switchover by using `Move-ActiveMailboxDatabase`. Listing 18-9 shows the syntax and usage.

LISTING 18-9 `Move-ActiveMailboxDatabase` cmdlet syntax and usage

Syntax

```
Move-ActiveMailboxDatabase -Identity DatabaseName
[-SkipClientExperienceChecks <$true | $false>] [-SkipHealthChecks
<$true | $false>] [-SkipLagChecks <$true | $false>] {AddtlParams}
```


Move-ActiveMailboxDatabase -Server **ServerName** {AddtlParams}

{AddtlParams}
[-ActivateOnServer **ServerOnWhichToActivate**] [-MountDialOverride
{"Lossless" | "GoodAvailability" | "BestAvailability"
"BestEffort" | "None"} [-DomainController **FullyQualifiedName**]
[-MoveComment **Comment**] [-SkipActiveCopyChecks <\$true | \$false>]
[-SkipClientExperienceChecks <\$true | \$false>]
[-SkipHealthChecks <\$true | \$false>] [-SkipLagChecks <\$true | \$false>]
[-TerminateOnWarning <\$true | \$false>]

Usage

Move-ActiveMailboxDatabase -Identity "Engineering Primary Database"
-ActivateOnServer "MailServer86" -MountDialOverride "Lossless"

The -MountDialOverride parameter of the Move-ActiveMailboxDatabase cmdlet controls the way databases mount on switchover or failover. Every Mailbox server has a database automount setting and the default is Best Availability. Values you can select to control the database mount behavior include:

- **None** The database uses the currently configured setting for automatically mounting.
- **Lossless** The database does not automatically mount until all logs that were generated on the original source server have been copied to the target node.
- **Good Availability** The database automatically mounts if the copy queue length is less than or equal to 6. If the queue length is greater than 6, Exchange attempts to replicate the remaining logs to the target server and mounts the databases once the queue length is less than or equal to 6.
- **Best Availability** The database automatically mounts if the copy queue length is less than or equal to 12. The copy queue length is the number of logs that need to be replicated. If the queue length is greater than 12, Exchange attempts to replicate the remaining logs to the target server and mounts the databases once the queue length is less than or equal to 12.
- **Best Effort** The database automatically mounts regardless of the length of the copy queue. As this option essentially forces the database to mount with any amount of log loss, I don't recommend using this value unless you are certain you want to accept what could be a large amount of data loss.

REAL WORLD You can set the default database automount setting for a Mailbox server by using the -AutoDatabaseMountDial parameter of the Set-MailboxServer cmdlet. If you specify either Best Availability or Good Availability and all of the data has not been replicated to the target server, you might lose some mailbox data; however, the transport dumpster feature (which is enabled by default) helps protect against data loss by resubmitting messages that are in the transport dumpster queue. Because of latency problems or other issues, specifying

one of these values can result in a database not being mounted, and you might need to use the `-AcceptDataLoss` parameter with `Mount-Database` to force the database to mount after a specified amount of time.

Checking Continuous Replication Status

You can use `Test-ReplicationHealth` to monitor continuous replication and determine the health and status of the underlying cluster service, quorum, and network components. By default, `Test-ReplicationHealth` performs the following tests:

- **ActiveManager** Verifies that the instance of Active Manager is running on the server.
- **ClusterNetwork** Verifies that all cluster-managed networks on the server are available.
- **ClusterService** Verifies that the Cluster service is running and reachable on the server.
- **DagMembersUp** Verifies that all DAG members are available, running, and reachable.
- **DBCopiedFailed** Checks whether any mailbox database copies are in a state of Failed on the server.
- **DBCopiedSuspended** Checks whether any mailbox database copies are in a state of Suspended on the server.
- **DBDisconnected** Checks whether any mailbox database copies are in a state of Disconnected on the server.
- **DBInitializing** Checks whether any mailbox database copies are in a state of Initializing on the server.
- **DBLogCopyKeepingUp** Verifies that log copying and inspection by the passive copies of databases on the server are able to keep up with log generation activity on the active copy.
- **DBLogReplayKeepingUp** Verifies that replay activities for the passive copies of databases on the server are able to keep up with log copying and inspection activity.
- **FileShareQuorum** Verifies that the witness server and witness directory and share configured for the DAG are reachable.
- **QuorumGroup** Verifies that the default cluster group (quorum group) is in a healthy and online state.
- **ReplayService** Verifies that the Microsoft Exchange Replication service is running and reachable on the server.
- **TasksRpcListener** Verifies that the tasks remote procedure call (RPC) server is running and reachable on the server.
- **TcpListener** Verifies that the TCP log copy listener is running and reachable on the server.

Listing 18-10 shows the syntax and usage for `Test-ReplicationHealth`. You can include monitoring events and performance counters in the results if desired. To do this, set the `-MonitoringContext` parameter to `$true`. Use `-OutputObjects` to output an array of results.

LISTING 18-10 Test-ReplicationHealth cmdlet syntax and usage

Syntax

```
Test-ReplicationHealth [-Identity MailboxServerToCheck]
[-ActiveDirectoryTimeout Timeout ] [-DomainController DCName]
[-MonitoringContext <$true | $false>] [-OutputObjects]
[-TransientEventSuppressionWindow Timeout]
```

Usage

```
Test-ReplicationHealth -Identity MailServer15 -ActiveDirectoryTimeout 30
-OutputObjects
```

Restoring Operations After a DAG Member Failure

If a Mailbox server has failed and cannot be recovered, you can recover operations in one of two ways:

- You can remove the configuration settings for the Mailbox server from the database availability group.
- You can install a new server and then restore the roles and settings for the original server.

Before you can remove the configuration settings for a Mailbox server, you'll need to remove any mailbox database copies that the server hosted. Use `Get-MailboxDatabaseCopyStatus` to list mailbox database copies and then use `Remove-MailboxDatabaseCopy` to remove the copies. Next, use the `Remove-DatabaseAvailabilityGroupServer` cmdlet to remove the configuration settings for the Mailbox server from the database availability group. After you remove the configuration settings, all settings associated with the Mailbox server are gone.

REAL WORLD Before you install a new server and then restore the roles and settings of the original server, you should confirm the install location for Exchange 2016 on the original server. If Exchange 2016 is installed in a location other than the default location, you can use the `/TargetDir` option during the setup of the new server to specify an install location. Otherwise, setup will use the default location for the installation. You can determine the install location for the original server by completing the following steps:

1. In `ADSIEdit.msc` or `LDP.exe`, navigate to `CN=ExServerName,CN=Servers,CN=First Administrative Group,CN=Administrative Groups,CN=ExOrg Name,CN=Microsoft Exchange,CN=Services,CN=Configuration,DC=DomainName,CN=Com`.
2. Right-click the Exchange server object for the failed server, and then select `Properties`.
3. The value of the `msExchInstallPath` attribute shows the installation path for the failed server.

To install a new server and then restore the roles and settings of the original server, complete the following steps:

1. Use `Get-MailboxDatabase` to list any replay lag or truncation lag settings for mailbox database copies that were hosted on the server being recovered. Enter the following command to list the databases associated with a specific server by their display name and lag settings:

```
Get-MailboxDatabase -server ServerName
```

Where *ServerName* is the name of the failed server. After you list the databases associated with the server, list the lag settings for each database in turn by entering the following command:

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

Where *DatabaseName* is the name of a database hosted on the failed server.

NOTE Alternatively, you can list all the databases associated with a specific server by their display name and lag settings by entering the following:

```
Get-MailboxDatabase -server Name | Get-MailboxDatabase |  
fl name, *lag*
```

Where *Name* is the name of the failed server. In this example, you examine the mailbox databases on MailServer24:

```
Get-MailboxDatabase -server MailServer24 | Get-MailboxDatabase  
| fl name, *lag*
```

2. After you list the databases associated with the server by name and lag times, you need to remove any mailbox database copies the server hosted by entering:

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

Where *DatabaseName* is the name of the database copy to remove and *ServerName* is the name of the failed server, such as:

```
Remove-MailboxDatabaseCopy EngDatabase\MailServer24
```

3. Remove the failed server's configuration from the DAG by entering the following command:

```
Remove-DatabaseAvailabilityGroupServer -Identity DagName  
-MailboxServer ServerName -ConfigurationOnly
```

Where *DagName* is the name of the DAG and *ServerName* is the name of the failed server, such as:

```
Remove-DatabaseAvailabilityGroupServer -Identity EastCampusDag1  
-MailboxServer MailServer24 -ConfigurationOnly
```

4. In Active Directory Users And Computers, locate and select the computer account for the failed server. On the Action menu, select Reset Account. When prompted to confirm, select Yes and then select OK.

5. Rename the new server so that it has the same name as the failed server. On the new server, run Exchange 2016 Setup with the /m:RecoverServer switch to have Setup read the failed server's configuration information from Active Directory. After Setup gathers the server's configuration information from Active Directory, Setup installs the original Exchange files and services on the server, restoring the roles and settings that were stored in Active Directory.
6. When the setup of the new server is complete, add the server to the DAG by entering the following command:

```
Add-DatabaseAvailabilityGroupServer -Identity DagName  
-MailboxServer ServerName
```

Where *DagName* is the name of the DAG and *ServerName* is the name of the failed server, such as:

```
Add-DatabaseAvailabilityGroupServer -Identity EastCampusDag1  
-MailboxServer MailServer24
```

7. Add mailbox database copies to the server by entering the following command for each database copy to add:

```
Add-MailboxDatabaseCopy -Identity DatabaseName -MailboxServer ServerName
```

Where *DatabaseName* is the name of the database copy to add and *ServerName* is the name of the server you are configuring, such as:

```
Add-MailboxDatabaseCopy -Identity EngDatabase  
-MailboxServer MailServer24
```

If any of the database copies had replay lag or truncation lag times greater than 0, you can set those lag times by using the `-ReplayLagTime` and `-TruncationLagTime` parameters.

8. After the database copies have been configured, you can check their status by entering the following command:

```
Get-MailboxDatabaseCopyStatus -Server ServerName
```

Where *ServerName* is the name of the server you are configuring, such as:

```
Get-MailboxDatabaseCopyStatus -Server MailServer24
```

The databases and their content indexes should have a healthy status.

9. Verify replication health for the server by entering the following command:

```
Test-ReplicationHealth -Identity ServerName
```

Chapter 19. Configuring Exchange Databases

Microsoft Exchange Server 2016 stores mailboxes and associated user data in mailbox databases. The information stored in a particular database isn't exclusive to mailboxes and their associated user data, however. Exchange Server maintains other related information within databases as well, including information about Exchange logons and mailbox usage. Exchange also maintains information about full-text indexing in mailbox databases, although the actual content indexes are stored in separate files.

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, "Creating and Managing Database Copies" later in this chapter. You create passive copies of databases as part of a high-availability configuration as discussed in Chapter 18, "Implementing Availability Groups."

Getting Started with Active Mailbox Databases

Every Mailbox server deployed 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 and multiple log files associated with it. These files are stored in a location that you specify when you create or modify the mailbox database.

Planning for Mailbox Databases

Within an Exchange organization, mailboxes are the normal delivery location for messages. Mailboxes 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 deploy 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 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 2016.
- **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 Office Outlook 2010 and later clients access the default OAB and default public folder hierarchy on your organization's Mailbox servers. Exchange 2016 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 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:

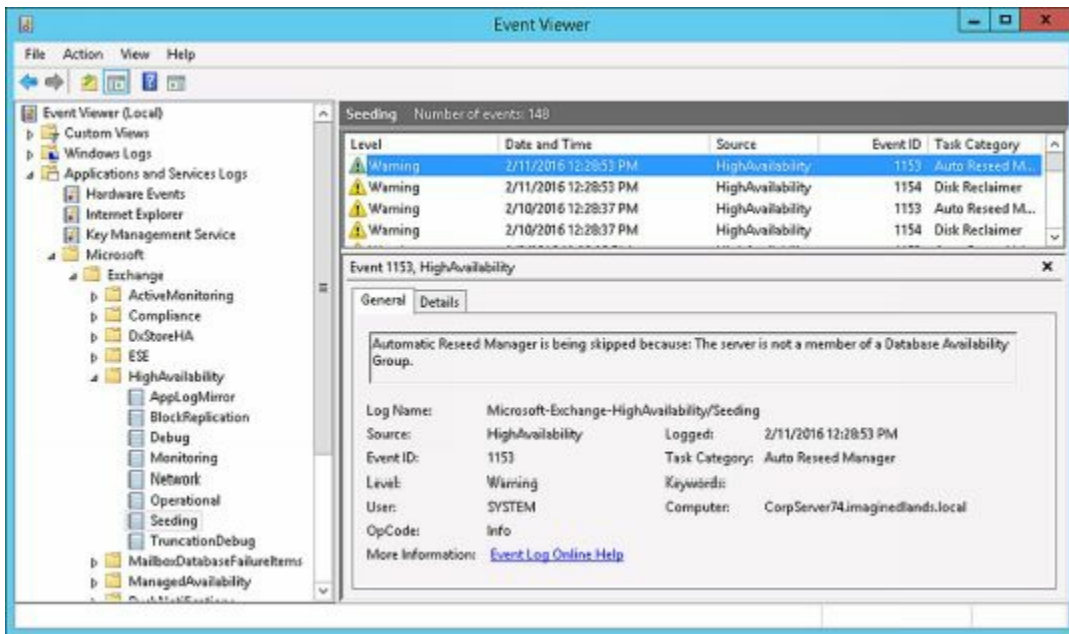
You have 5 volumes mounted under C:\PrimaryVols as C:\PrimaryVols\Volume1, C:\PrimaryVols\Volume2, C:\PrimaryVols\Volume3, C:\PrimaryVols\Volume4, and C:\PrimaryVols\Volume5.

You create 3 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 2 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.



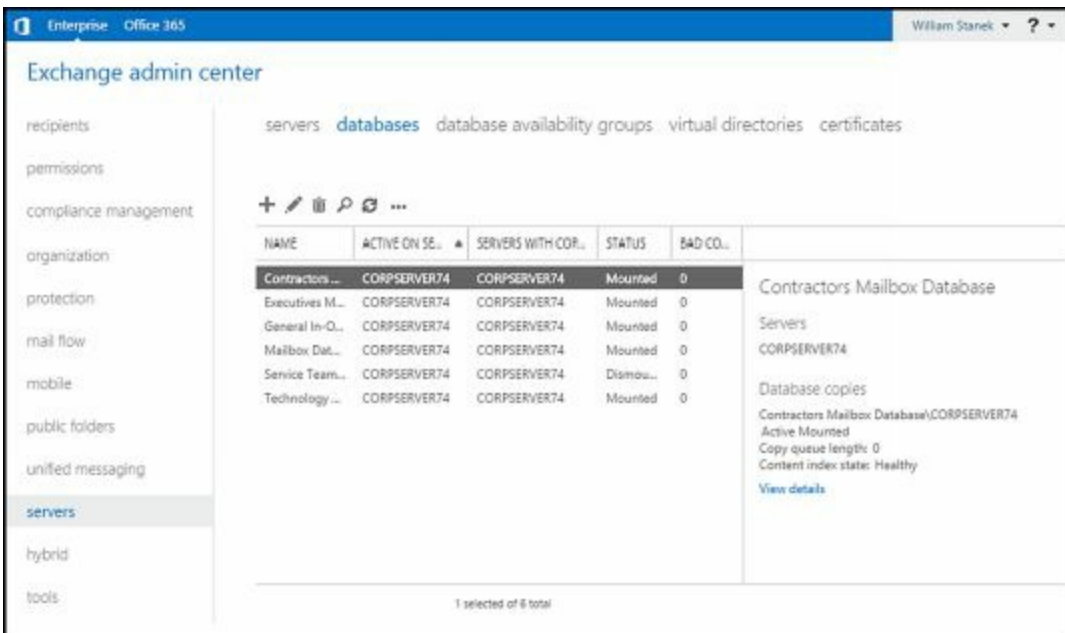
Creating and Managing Active Databases

In the Exchange Admin Center, select Servers in the Navigation menu, and then select Databases to view the currently available databases in the organization. Database are listed by name, active server, servers with copies, status and bad copy count.

Creating Mailbox Databases

You use the New Mailbox Database wizard to create mailbox databases. Each database has an associated database file path for its .edb file and an associated log folder for its logs. Any new mailbox databases you create using the Exchange Admin Center are configured to use the mailbox provisioning load balancer by default. If you create databases using the 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.



You can create a mailbox database by completing the following steps:

1. In the Exchange Admin Center, select Servers in the Navigation menu, and then select Databases. In the main pane, you should see a list of active databases that are available in the Exchange organization.

new database


*Mailbox database

*Server

Database file path:

Log folder path:

Mount this database

2. Click New () to open the New Database dialog box.
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. Click Browse to the right of the Server text box to open the Select Server dialog box. Mailbox servers are listed by name, version, and exact build as well as associated database availability group, if applicable.


NAME	VERSION	DATABASE AVAILABILITY GROUP
MAILSERVER21	Version 15.1 (Bu...	WestCampusDAG1
MAILSERVER42	Version 15.1 (Buil...	WestCampusDAG1

5. Select the Mailbox server that will host the mailbox database, and then 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.

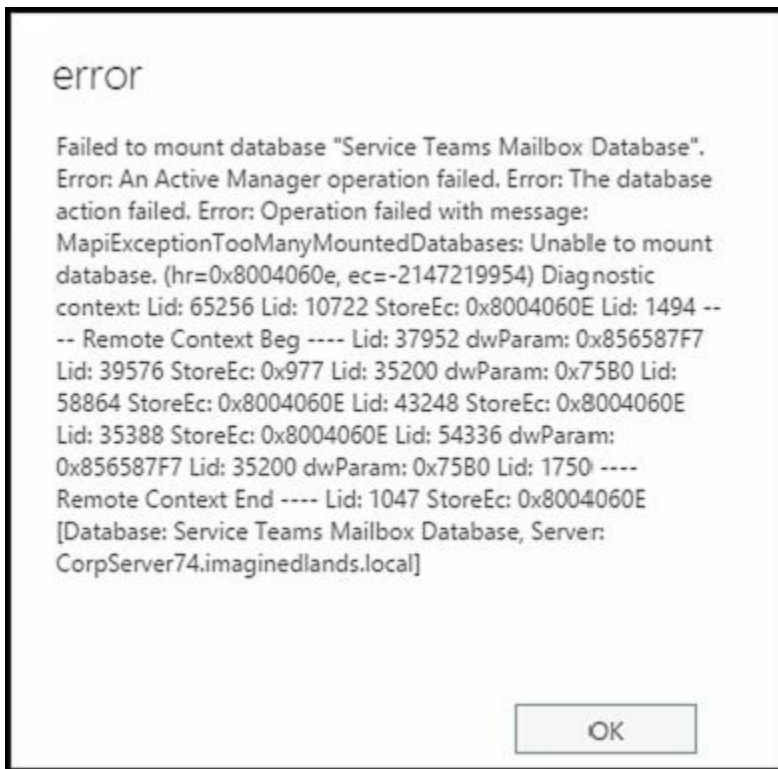
REAL WORLD Exchange creates folders if they do not exist, which is a good thing except when you mistype the intended path. Rather than type in a long file path, you might want to use copy and paste. In Windows Explorer, navigate to the exact folder path you want to use. Click in the folder path on the 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. Click Save to create the mailbox database, and then 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.

NOTE In Exchange Admin Center, you may need to click Refresh () to see the newly created database under Servers > Databases.

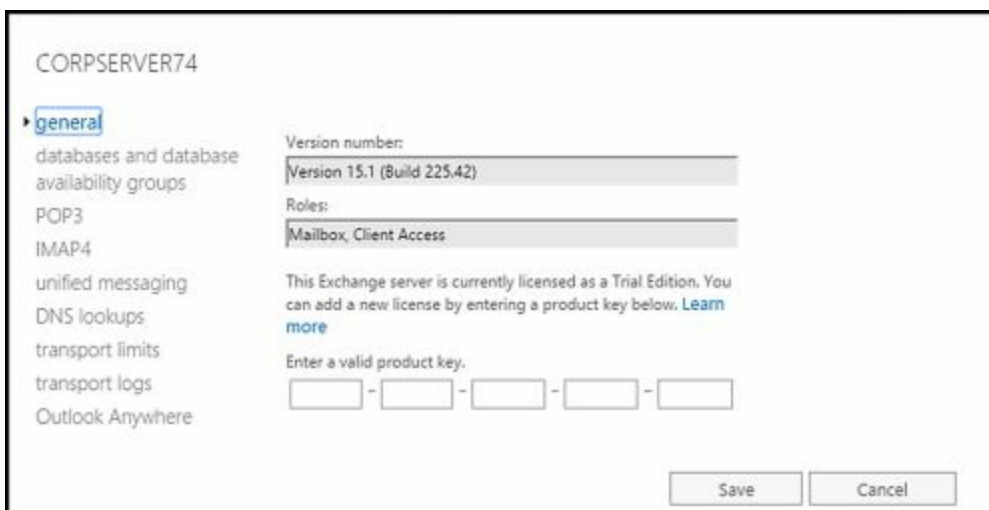
Exchange Server 2016 Standard edition supports up to five databases. Exchange Server 2016 Enterprise edition supports up to 100 databases. However, if you install Exchange Server 2016 Enterprise edition but forget to enter the product key, the server runs in Trial mode and only supports up to five databases as well.

Exchange Server can't mount more databases than are supported by the edition you are using. When you create more databases than are supported, Exchange will be unable to mount the database. In the error details, you'll also see a message stating MapiExceptionTooManyMountedDatabases: Unable to mount database.



You can resolve the too many databases problem by reducing the number of databases on the server or simply creating the database on a different server. If the server is running a standard or trial edition of Exchange Server 2016, you can upgrade the server to Enterprise edition to resolve the problem by completing these steps:

1. In Exchange Admin Center, select Servers on the Navigation menu, and then select Servers.
2. 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.



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. Click OK. For the change to take effect, you must restart the Microsoft Exchange

Information Store service.

In Exchange Management Shell, you can create mailbox databases by using the `New-MailboxDatabase` cmdlet. Listing 19-1 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 19-1 `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]
```

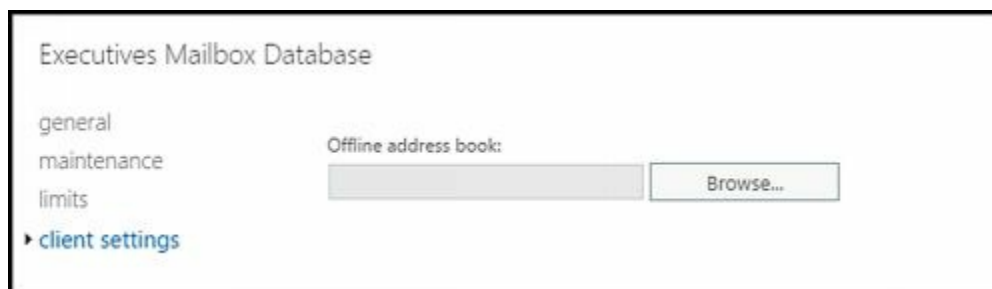
Usage

```
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-click the database you want to configure.
2. In the Properties dialog box, 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. Click Browse, select the OAB you want to use, and then click OK. Click Save to apply the changes.

In Exchange Management Shell, you can set the default OAB for mailbox databases by using the Set-MailboxDatabase cmdlet. Listing 19-2 provides the syntax and usage.

LISTING 19-2 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 seven days. In most cases, you'll want deleted messages to be maintained for a minimum of five to seven days and deleted mailboxes to be maintained for a minimum of three to four weeks. An interval of five to seven 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-click the database you want to configure.
 2. In the Properties dialog box, on the Limits page, 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.



The screenshot shows the 'Limits' page for a mailbox database named 'Technology Mailbox Database'. The page is divided into sections: 'general', 'maintenance', 'limits', and 'client settings'. The 'limits' section is active and contains the following settings:

- *Issue a warning at (GB): 1.9
- *Prohibit send at (GB): 2
- *Prohibit send and receive at (GB): 2.3
- *Keep deleted items for (days): 14
- *Keep deleted mailboxes for (days): 30
- Don't permanently delete items until the database is backed up

At the bottom of the page, there are 'Save' and 'Cancel' buttons.

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 nondelivery 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.
 - **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 seven 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.
- Times that are used for quota notification are filled in with a dark bar.
 - Times that aren't used for quota notification are blank.

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 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.

Show the time in hours
 Show the time in 15-minute intervals

	Midnight (AM)											Noon (PM)												
All	12	1	2	3	4	5	6	7	8	9	10	11	12	1	2	3	4	5	6	7	8	9	10	11
Sunday																								
Monday																								
Tuesday																								
Wednesday																								
Thursday																								
Friday																								
Saturday																								

Quota notification hours
 Non-quota notification hours

OK Cancel

5. If you customized the notification schedule, click OK to close the Customize Quota Notification Schedule dialog box.
6. Click Save to apply your settings.

In Exchange Management Shell, you can set limits for mailbox databases by using the Set-MailboxDatabase cmdlet. Listing 19-3 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 2016.

LISTING 19-3 Using the Set-MailboxDatabase cmdlet to set limits

Syntax

```
Set-MailboxDatabase [-Identity MailboxDatabase]  
[-AllowFileRestore <$true | $false>] [-BackgroundDatabaseMaintenance <$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.

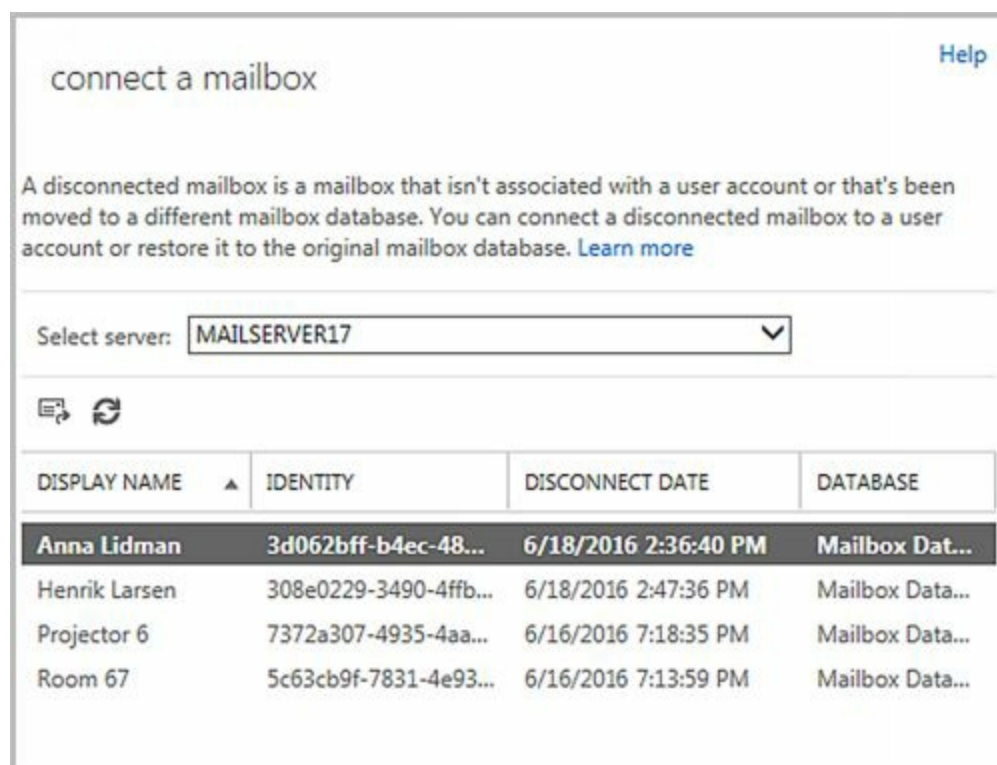
When you move mailboxes between databases, mailboxes in the original (source) database are soft deleted. This means they are disconnected, marked as soft deleted, but

retained in the original database until the deleted mailbox retention period expires. In Exchange Management Shell, you can use a DisconnectReason of “SoftDeleted” to find soft-deleted mailboxes.

@techjob

To recover a deleted mailbox, complete the following steps:

1. In the Exchange Admin Center, select Recipients in the Navigation menu, and then select Mailboxes.
2. Click the More button (**⋮**), 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.



connect a mailbox Help

A disconnected mailbox is a mailbox that isn't associated with a user account or that's been moved to a different mailbox database. You can connect a disconnected mailbox to a user account or restore it to the original mailbox database. [Learn more](#)

Select server:

⌵ ↻

DISPLAY NAME ▲	IDENTITY	DISCONNECT DATE	DATABASE
Anna Lidman	3d062bff-b4ec-48...	6/18/2016 2:36:40 PM	Mailbox Dat...
Henrik Larsen	308e0229-3490-4ffb...	6/18/2016 2:47:36 PM	Mailbox Data...
Projector 6	7372a307-4935-4aa...	6/16/2016 7:18:35 PM	Mailbox Data...
Room 67	5c63cb9f-7831-4e93...	6/16/2016 7:13:59 PM	Mailbox Data...

3. In the Connect A Mailbox dialog box, use the selection list provided to select the server in which you want to look for disconnected mailboxes.
4. Click the mailbox to restore it, and then 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 19-4.

LISTING 19-4 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 "Don Harmon"
-Database "Accounting Mail" -User "TVPRESS\donh" -Alias "donh"
```

```
Connect-Mailbox -Identity "Don Harmon"
-Database "Accounting Mail" -LinkedDomainController CorpServer72
-LinkedMasterAccount "TVPRESS\donh"
```

Recovering Deleted Items from Servers

You can recover deleted items from mailbox servers 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 higher to recover deleted items from a Mailbox server, complete the following steps:

1. [Log on as the user who deleted the message, and then start Outlook.](#)

2. In the email folder list, select Deleted Items.
3. With Home selected on the navigation pane, select Recover Deleted Items From Server. If this option isn't available, make sure you've access the Exchange account in Outlook and are using online mode.
4. The Recover Deleted Items From dialog box appears. Select the items you want to recover, and then click the Recover Selected Items button.
5. Items you've recovered are copied to the Deleted Items folder. In the left pane, click Deleted Items.
6. In the Deleted Items folder, right-click items you want to keep, select Move, and then click Other Folder.
7. In the Move Items dialog box, select the folder to which the item should be moved, such as Inbox, and then click OK.

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 2016 and *yourdomain.com* is a placeholder for your external domain name, such as <https://mail.tvpress.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 **imaginedlands\bertk**, or *user@domain* format, such as **bertk@imaginedlands.com** . Type the password, and then click Sign In.
3. In the left pane, 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 checkbox for items you want to recover.
5. Click the Recover button, and then click OK. Items you select will be restored to their default folders.

Creating and Managing Database Copies

Mailbox databases are either active or passive. 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 18, “Implementing Availability Groups.” 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.

Creating Mailbox Database Copies

Once 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 2016 mailbox database to a server running an earlier version of Exchange.

A 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.

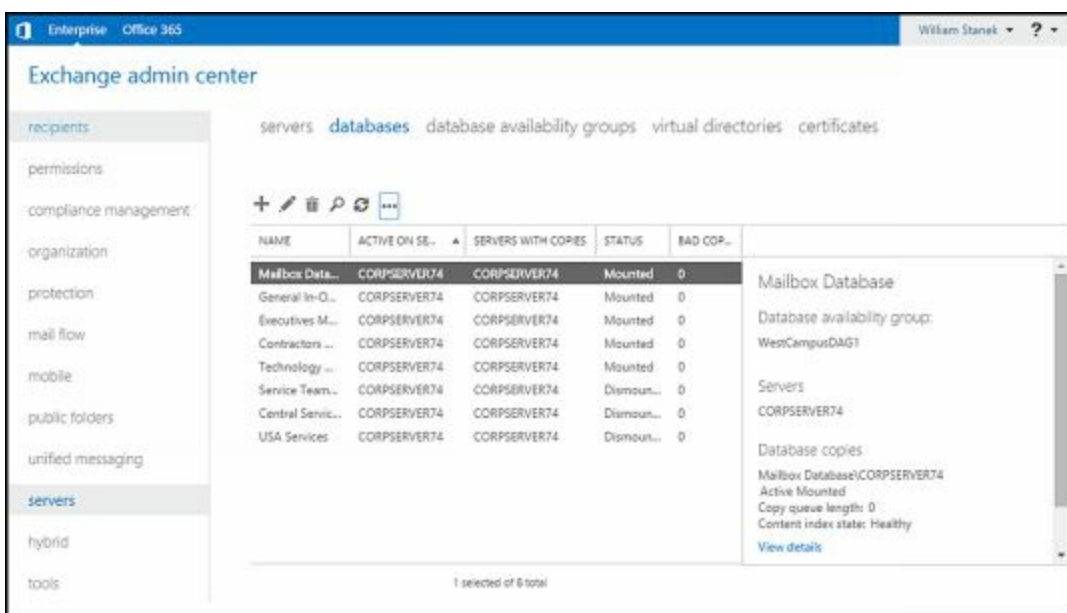
Because all copies of a database use the same path on each server containing a copy, the database and log file paths for a database copy on each Mailbox server must not conflict with any other database paths. You need to ensure the database and log file paths for the database copy can be created in the 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 Navigation menu, 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. Click the More button (**⋮**), and then select Add A Database Copy. This opens the Add Mailbox Database Copy dialog box, which lists the servers that 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.

4. Click Browse. Select the Mailbox server that will host the mailbox database copy, and then 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 2016 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 “Active Manager Framework” in Chapter 18.

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. 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 Exchange Management Shell, you can create mailbox database copies by using the `Add-MailboxDatabaseCopy` cmdlet. Listing 19-5 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.

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.

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.

LISTING 19-5 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
```

Configuring Database Copies

Database copies have associated replay, truncation, and preference values. 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 Navigation menu, 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. 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 click Save.

In Exchange Management Shell, you can set replay, truncation, and preference values for mailbox database copies by using the Set-MailboxDatabaseCopy cmdlet. Listing 19-6 provides the syntax and usage.

LISTING 19-6 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 2016, lagged copies automatically play down log files as necessary to accommodate adverse conditions, such as the following:

- If Exchange 2016 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 2016 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 2016 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
```

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

You specify the low disk space 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 2016 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

You may need to suspend replication for a database copy as part of planned maintenance or for other reasons.. 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 Navigation menu, 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. Click Save to suspend continuous replication.

To resume replication later, click the Resume option in the details pane. If a suspend comment was provided, you can read the comment. Then click Yes to resume continuous replication.

In Exchange Management Shell, you can suspend and resume replication by using `Suspend-MailboxDatabaseCopy` and `Resume-MailboxDatabaseCopy`, respectively. Listings 19-7 and 19-8 provide the syntax and usage. If you use the `-ActivationOnly` parameter to suspend activation only, the database cannot be activated until you resume replication without specifying the `-ReplicationOnly` parameter. You can use the `-ReplicationOnly` parameter to resume replication without affecting the activation setting. For example, if the `-ActivationSuspended` parameter was set to `$True`, the parameter remains set to `$True`.

LISTING 19-7 `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 19-8 `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
```

Where *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\MailServer96"  
-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 Navigation menu, and then select Databases. Next, select the database copy you activated. In the Details pane, click View Details.
4. Once 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
```

Where *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  
MailServer96 -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 Navigation menu, 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

Where 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 Navigation menu, 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 FailedAndSuspended 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. Once 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:

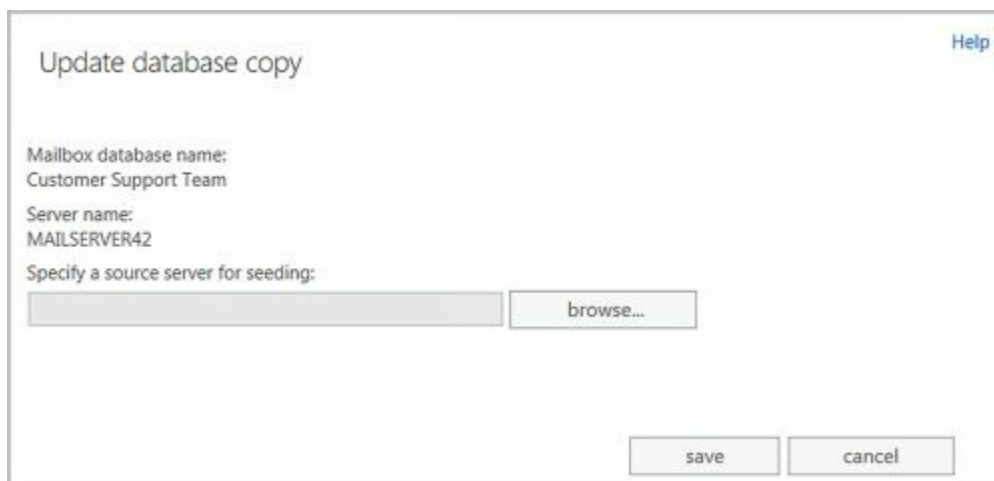
- When you seed a database using the Exchange Admin Center, both the database copy and the content index catalog are seeded. In Exchange Management Shell, you can specify that only the database copy should be seeded using the `-DatabaseOnly` parameter or that only the context index catalog should be 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 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 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 Exchange Management Shell, use `-NetworkCompressionOverride` and `-NetworkEncryptionOverride` to override the network compression and encryption settings, respectively.

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

1. In the Exchange Admin Center, select Servers in the Navigation menu, 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.

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.



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, click Browse. In the Select Mailbox Server dialog box, select the source server hosting the passive copy you want to use and then click OK.
5. Click Save to begin seeding. If an error occurred, you need to take the appropriate corrective action. Click Close.

To seed a database copy in Exchange Management Shell, you use the Update-MailboxDatabaseCopy cmdlet. Listing 19-9 provides the syntax and usage. You can use the `-Force` parameter when seeding programmatically, and you will not be prompted for administrative input.

LISTING 19-9 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

Monitoring the health and status of database copies is important 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 Navigation menu, 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. 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.

When you are working with database copies, the Copy Status shows the current health and status of replication for the database copy. The possible status values as well as any corrective action that might be required include:

- **Activation Suspended** The copy has been manually blocked from activation by an

administrator. To correct this status, allow activation, if appropriate.

- **Disconnected And Healthy** The copy 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** The copy 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** The copy is offline and not accepting client connections. Applies only to the active mailbox database. To correct this status, mount the database if maintenance is complete.
- **Dismounting** The copy is going offline and terminating client connections. Applies only to the active mailbox database.
- **Failed** The copy 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** The copy 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.
- **Healthy** The copy is successfully copying and replaying log files, or has successfully copied and replayed all available log files.
- **Initializing** The copy 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** The copy is online and accepting client connections. Applies only to active mailbox database.
- **Mounting** The copy is coming online and not yet accepting client connections. Applies only to active mailbox database.
- **Resynchronizing** The copy 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** The copy 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** The copy cannot connect to the replication service. To correct this status, start or restart the Microsoft Exchange Replication service on the server that hosts the mailbox database copy.
- **Single Page Restore** The copy had a single page error, and this error is being corrected automatically.

- **Suspended** The copy is in a suspended state as a result of an administrator manually suspending the database copy. To correct this status, resume replication if appropriate

In Exchange Management Shell, you can check the health and status of replication by using the `Get-MailboxDatabaseCopyStatus` cmdlet. Listing 19-10 provides the syntax and usage.

LISTING 19-10 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"
```

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 Navigation menu, 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. Click the Remove option for the database copy you want to remove.
4. When prompted to confirm, click Yes. If an error occurred, you need to take the appropriate corrective action. Click Close.

In Exchange Management Shell, you can remove a database copy by using the `Remove-MailboxDatabaseCopy` cmdlet. Listing 19-11 provides the syntax and usage.

LISTING 19-11 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}
```


Maintaining Mailbox Databases

After exploring how to create and use databases, let's look at some general techniques you'll use to manage databases. Keep in mind that these techniques apply only to active mailbox databases. Passive copies of mailbox databases are managed as discussed in the "Creating and Managing Database Copies" section earlier in this chapter.

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 2016 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 2016 will be unable to mount the database. Exchange Server 2016 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 2016 is unable to find the file on the disk. When Exchange Server 2016 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 2016 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 2016 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 2016 displays a warning before mounting the database and recreating the database file. 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.

Checking Database Status

Mailbox databases have several associated states, including

- 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 Navigation menu, 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 Exchange Management Shell, you can determine the status of all databases or specific databases by using the Get-MailboxDatabase. Listing 19-12 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 19-12 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

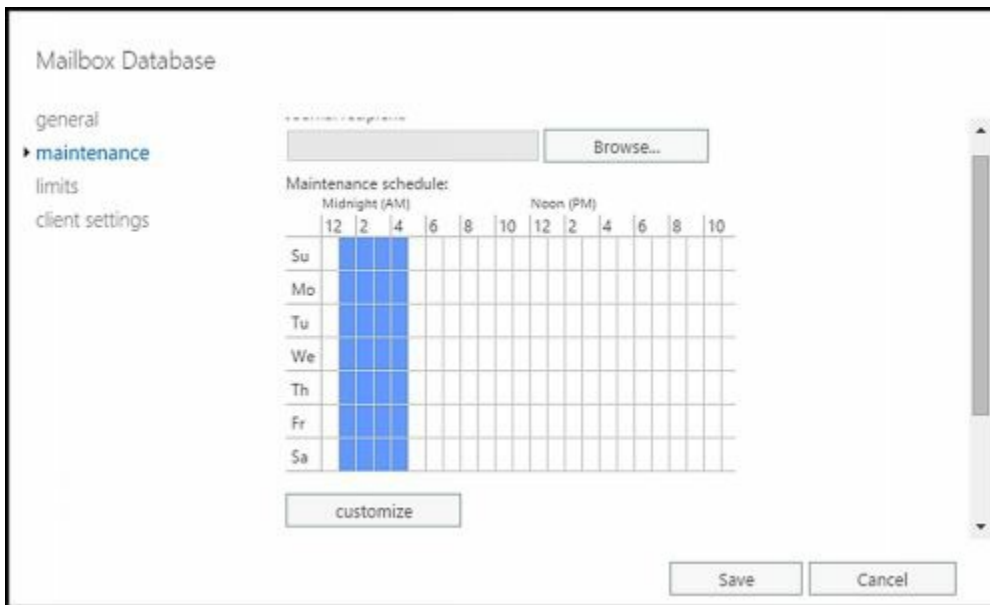
Get-MailboxDatabase –Status | format-table Name,
Mounted, BackupInProgress, BackgroundDatabaseMaintenance

Setting the Maintenance Interval

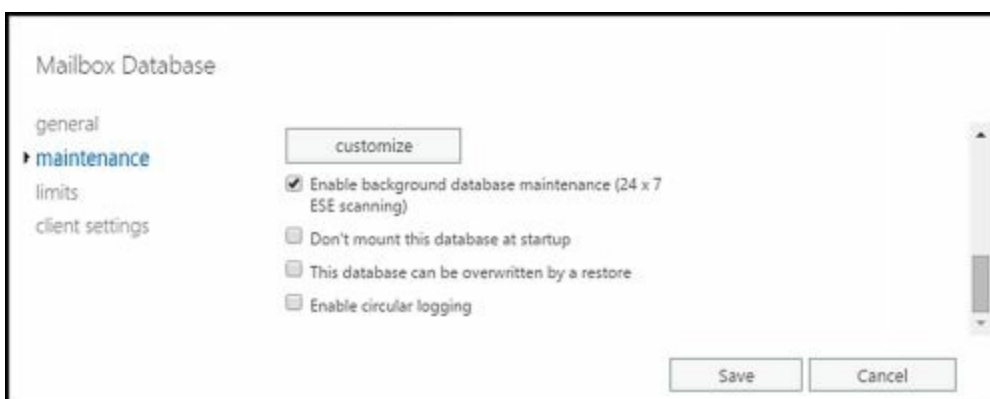
You should run maintenance routines against databases on a daily basis. The maintenance routines organize the databases, clear out extra space, and perform other essential housekeeping tasks. By default, the automatic background maintenance does some of this work, and Exchange Server runs extended, foreground 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 Navigation menu, 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-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.
- Times that are used for maintenance are filled in with a dark bar.
 - Times that aren't used for maintenance are blank.

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 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. 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, 7 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. Click OK.



In Exchange Management Shell, you can configure the maintenance schedule for a database by using `Set-MailboxDatabase`. Listing 19-13 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 19-13 Setting the maintenance schedule

Syntax

`Set-MailboxDatabase -Identity DatabaseIdentity`

```
[-MaintenanceSchedule Schedule]  
[-BackgroundDatabaseMaintenance <$true | $false>]
```

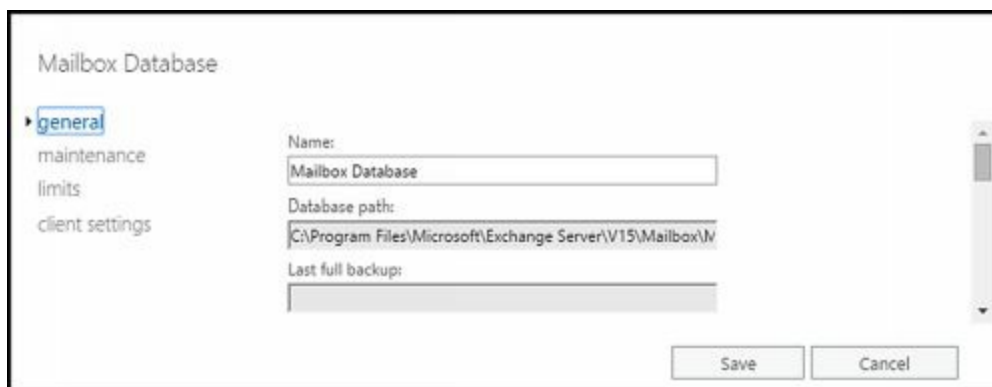
Usage

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

Renaming Databases

To rename a database, follow these steps:

1. In the Exchange Admin Center, double-click the database with which you want to work. This opens a properties dialog box for the database. The General page is selected by default.
2. In the Name text box, type the new name for the database. 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 Exchange Management Shell, you can rename databases by using the `-Name` parameter of the `Set-MailboxDatabase`. Listing 19-14 provides the syntax and usage.

LISTING 19-14 Renaming a database

Syntax

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

Usage

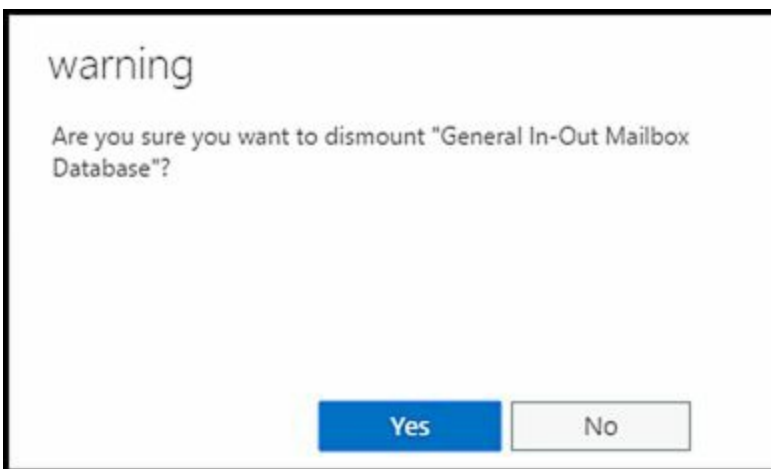
```
Set-MailboxDatabase -Identity "Eng DB"  
-Name "Engineering Mail Database"
```

Mounting and Dismounting 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 Navigation menu, 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. Click the More button (**⋮**), and then select Dismount. When prompted, confirm the action by 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, clicking the More button, and then selecting Mount. When prompted, confirm the action by clicking Yes.

In Exchange Management Shell, you can dismount and mount databases by using the Dismount-Database and Mount-Database cmdlets, respectively. Listing 19-15 provides the syntax and usage for these cmdlets.

LISTING 19-15 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

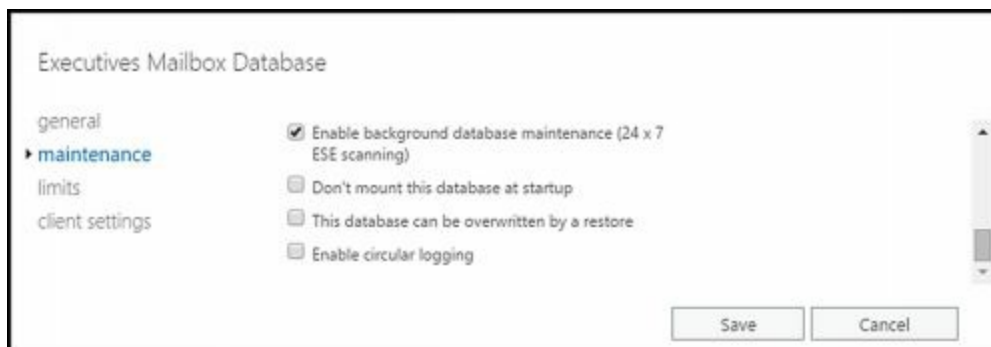
Mount-Database –Identity "Eng DB"

Configuring Automatic Mounting

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 Navigation menu, 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-click the database with which you want to work.
3. On the Maintenance page, do one of the following and then click Save:
 - To ensure that a database isn't mounted on startup, select the Don't Mount This Database At Startup check box.
 - To mount the database on startup, clear the Don't Mount This Database At Startup check box.



In Exchange Management Shell, you can enable or disable automatic mounting at startup by using the Set-MailboxDatabase. Listing 19-16 provides the syntax and usage for controlling automatic mounting.

LISTING 19-16 Controlling automatic mounting

Syntax

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

Usage

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

Moving Databases

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
```

Where *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 Exchange Management Shell, you can move databases by using the *Move-DatabasePath* cmdlet. Listing 19-17 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 19-17 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*
```

Where *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
```

Where *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. As 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 MExchangeSearch  
Start-Service MExchangeSearch
```

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-MailboxDatabaseCopy 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]
```

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

Deleting Databases

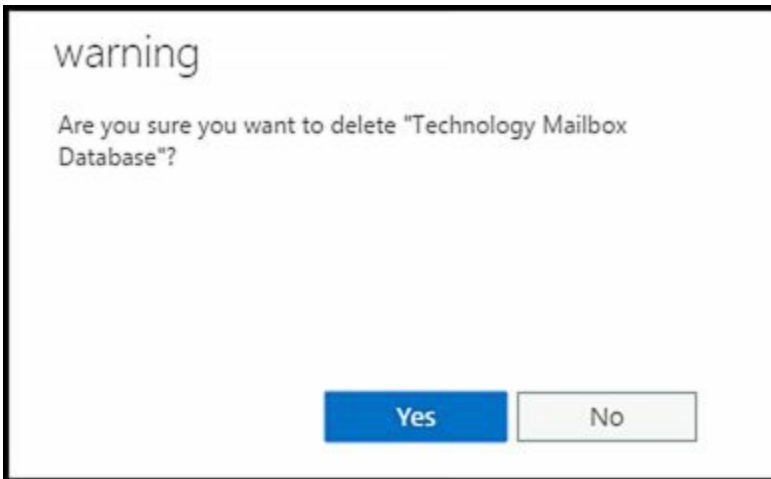
Before deleting a mailbox database, you must disable, move or remove all mailboxes, archive mailboxes, public folder mailboxes and arbitration mailboxes from the database. To help you with this process, use the following techniques:

- Get a list of all mailboxes in the database by running the command **Get-MailboxDatabase -Database *DatabaseName* | Get-Mailbox**. These are the user mailboxes that must be moved or removed.
- Get a list of all archive mailboxes in the database by running the command **Get-MailboxDatabase -Database *DatabaseName* | Get-Mailbox -Archive**. These are the archive mailboxes that must be moved or removed.
- Get a list of all arbitration mailboxes in the database by running the command **Get-MailboxDatabase -Database *DatabaseName* | Get-Mailbox -Arbitration**. These are the arbitration mailboxes that must be moved or removed.
- Disable a non-arbitration mailbox so that you can delete the mailbox database by running the command **Disable-Mailbox *MailboxId***.
- Disable an archive mailbox so you can delete the mailbox database by running the command **Disable-Mailbox *MailboxId* -Archive**.
- Disable a public folder mailbox so that you can delete the mailbox database by running the command **Disable-Mailbox *MailboxId* -PublicFolder**.
- Rather than removing arbitration mailboxes, you should move them to another database using New-MoveRequest. If this is the last database in the organization, disable the arbitration mailbox instead by running the command **Disable-Mailbox**

ArMailboxID -Arbitration -DisableLastArbitrationMailboxAllowed.

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 clicking Yes. If the database contains any mailboxes, you'll see an error and will need to disable, move or remove the mailboxes before you can remove the database.



3. After removing the database, you need to delete any database and transaction log files from the server.

In Exchange Management Shell, you can delete databases by using the Remove-MailboxDatabase. Listing 19-18 provides the syntax and usage.

LISTING 19-18 Removing databases

Syntax

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

Usage

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

Managing Content Indexing

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.

Indexing Essentials

Content indexing enables fast searches and lookups through server-stored mailboxes. Exchange Server 2016 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 Office Outlook 2010, all users need to do is click in the Search box or press Ctrl+E to access the Search tools, click Search Options, and then click Advanced Find. In the Advanced Find dialog box, users can enter their search parameters and then 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).

Maintaining Exchange Store Search

Exchange Server 2016 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*DatabaseName* \ *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 2016 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 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 Exchange Management Shell follows:

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

```
Start-Service MExchangeFastSearch -ComputerName MailServer11
```

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

Resolving Indexing Issues

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
```

Where 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-ExchangeSearch. 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.

Chapter 20. Managing SMTP Connectors

SMTP connectors, Active Directory sites, and Active Directory links all have important roles to play in determining how Exchange routes and delivers messages in your organization. You can work with connectors, sites, and links in a variety of ways, but first you need to have a strong understanding about how connectors are used.

Send and Receive Connectors: The Essentials

In Exchange Server 2016, SMTP connectors are used to logically represent the connection between a source server and a destination server. Only transport servers have SMTP connectors and in an Exchange 2016 organization there are two primary types of transport servers:

- Mailbox servers
- Edge Transport servers

NOTE As Exchange 2016 can interoperate with Exchange 2013 and Exchange 2010, Exchange 2013 Mailbox servers and Exchange 2010 Hub Transport servers can also act as transports. Thus, in a mixed environment the transport servers can include Exchange 2010 Hub Transports as well as Mailbox servers and Edge Transports running Exchange 2013 and Exchange 2016.

How you configure an SMTP connector determines how Exchange Server transports messages using that connection. Because each SMTP connector represents a one-way connection, Exchange Server uses both Send and Receive connectors.

Understanding Send and Receive Connectors

Send connectors are logical gateways through which Exchange servers send all outgoing messages. When you create a Send connector, it is stored in Active Directory Domain Services or in Active Directory Lightweight Directory Services as a connector object.

Send connectors are not scoped to a single server; in fact, multiple servers can use a single Send connector for sending messages. Send connectors deliver mail by looking up a mail exchanger (MX) record on a DNS server, by looking up an Address (A) record, or by using a smart host as a destination. With DNS records, the DNS server settings you configure on the transport server are used for name resolution. You can configure different settings for internal and external DNS lookups if necessary. See the “Configuring Send Connector DNS Lookups” section of this chapter.

Receive connectors are logical gateways through which all incoming messages are received. When you create a Receive connector, it is stored in AD DS or in AD LDS as a connector object. Unlike Send connectors, Receive connectors are scoped to a single server and determine how that server listens for connections.

The permissions on a Receive connector determine from which other servers the connector will accept connections. The authentication mechanisms you configure for a Receive connector determine whether anonymous connections are allowed and the types of authentication that are permitted.

When you install your Mailbox servers, Exchange 2016 creates the connectors required for mail flow within your organization; however, to send mail outside your domain, you must create a Send connector to send mail to the Internet.

If your organization also uses Edge Transport servers, Exchange creates the additional Send and Receive connectors during the Edge Subscription process. You also can explicitly create Send and Receive connectors or automatically compute them from the organization topology by using Active Directory sites and site-link information.

REAL WORLD To enhance security and prevent malicious users from trying to determine internal infrastructure components, Exchange 2016 applies the header firewall feature to remove X-headers and routing headers from inbound and outbound messages automatically. *X-headers* are user-defined, unofficial headers added to messages during processing, filtering, or virus/spam checking that detail how a message was processed, filtered, or checked. *Routing headers* are standard SMTP headers that provide information about the various messaging servers used to deliver a message.

The exact headers removed from a message depend on the connector type. Receive connectors with the Internal type remove organization and forest X-headers from messages. Receive connectors with the Custom type remove routing headers (as long as permissions groups are not assigned). Internal or Partner type Send connectors remove organization and forest X-headers from messages. Custom type Send connectors remove routing headers (as long as permissions groups are not assigned).

Routing Messages within Sites

Exchange 2016 determines the ultimate destination for a message when routing messages and then uses a least-cost method to determine how to route the message. By default, Mailbox servers use Active Directory sites and the costs that are assigned to the Active Directory Internet Protocol (IP) site links to determine the least-cost routing path to other Mailbox servers in the organization. You can also specify an Exchange cost for links.

After a Mailbox server determines the least-cost routing path, the server routes messages over the link or links in this path, and in this way, a source Mailbox server relays messages to target Mailbox servers. By default, when there are multiple Active Directory sites between the source and destination server, the Mailbox servers along the path between the source server and the target server don't process or relay the messages in any way—with the following exceptions:

- If you want messages to be processed en route, you can configure an Active Directory site as a hub site so that Exchange routes messages to the hub site to be processed by the site's Mailbox servers before being relayed to a target server. The hub site must exist along the least-cost routing path between source and destination Mailbox servers.
- If a message cannot be delivered to the target site, the Mailbox server in the closest reachable site along the least-cost routing path of the target site queues the message for relay. The message is then relayed when the destination Mailbox server becomes available.

Sometimes during routing, messages must pass through transport servers in a hub site that isn't the ultimate destination but is part of the least-cost routing path. All transport servers in a hub site are considered part of the delivery group for that hop and a transport server is randomly selected to handle the message. As a message passes through a hub site, the randomly selected transport server queues and processes the message, routing the message along the least-cost path.

In cases in which a subscribed Edge Transport server is accessible only from the Active Directory site to which it is subscribed, messages must pass through a specific hub site to get to an Edge Transport server to be routed to the Internet or across premises. This happens because a subscribed Edge Transport server is only accessible from the Active Directory site to which it is subscribed.

Each routing destination has a delivery group that handles its message delivery. As discussed in Chapter 17, "Managing Exchange Organizations," in the "Front End Transport" section, an Active Directory site can be a delivery group but so can a routable DAG, a Mailbox delivery group, a group of connector source servers, or a list of expansion servers for dynamic distribution groups. When a DAG is the delivery group, the DAG itself is a routing boundary and the mailbox databases in the DAG are the routing destinations serviced by the related delivery group. Thus, a message can be sent from a mailbox on any transport server in the DAG to a mailbox on that server or on any other server in the DAG directly. As site boundaries don't apply, the member servers can be in different sites as well.

When a site is the delivery group, Exchange 2016 can use delayed fan-out to reduce the number of message transmissions by identifying recipients that share part of the routing path. On the other hand, when a site isn't the delivery group, Exchange 2016 selects one site in the destination delivery group with the least-cost routing path as the primary site. If multiple least-cost routing paths are available, the path with the fewest number of hops is chosen. If multiple paths are still available, the site nearest the destination is chosen based on the site name. Specifically, Exchange 2016 selects the site with the lowest alphanumeric sort order. For example, Seattle Site 1 is chosen before Seattle Site 2 and Alpha Site is chosen before Beta Site.

To display the configuration details of an Active Directory site, you can use the Get-AdSite cmdlet. If you don't provide an identity with this cmdlet, configuration information for all Active Directory sites is displayed.

Listing 20-1 provides the syntax and usage, as well as sample output, for the Get-AdSite cmdlet. Note that the output specifies whether the site is enabled as a hub site.

LISTING 20-1 Get-AdSite cmdlet syntax and usage

Syntax

```
Get-AdSite [-Identity 'SiteIdentity'] [-DomainController 'DCName']
```

Usage

```
Get-AdSite -Identity 'First-Seattle-Site' | fl
```

Output

```
RunspaceId      :
HubSiteEnabled  : False
InboundMailEnabled : True
PartnerId       : -1
MinorPartnerId  : -1
ResponsibleForSites : {}
Name            : First-Seattle-Site-Name
AdminDisplayName :
ExchangeVersion : 0.0 (6.5.6500.0)
DistinguishedName : CN=First-Seattle-Site-Name,CN=Sites,CN=Configuration
,DC=pocket-consultant,DC=com
Identity        : imaginedlands.com/Configuration/Sites/
First-Seattle-Site-Name
Guid            :
ObjectCategory   : imaginedlands.com/Configuration/Schema/Site
ObjectClass      : {top, site}
WhenChanged      : 2/15/2016 8:55:33 PM
WhenCreated      : 2/15/2016 8:55:33 PM
WhenChangedUTC   : 2/16/2016 4:55:33 AM
WhenCreatedUTC   : 2/16/2016 4:55:33 AM
OrganizationId   :
OriginatingServer : CorpServer27.imaginedlands.com
IsValid          : True
ObjectState      : Unchanged
```

To configure an Active Directory site as a hub site and override the default message routing behavior, you can use the `Set-AdSite` cmdlet with the `-HubSiteEnabled` parameter. To enable a site as a hub site, set the `-HubSiteEnabled` parameter to `$true`. To disable a site as a hub site, set the `-HubSiteEnabled` parameter to `$false`. You must have Enterprise Administrator rights to use the `-Name` parameter to change a site's name.

Listing 20-2 provides the syntax and usage, for the `Set-AdSite` cmdlet. Keep in mind that when a hub site exists along the least-cost routing path between source and destination Mailbox servers, messages are routed to the hub site for processing before they are relayed to the destination server.

LISTING 20-2 Set-AdSite cmdlet syntax and usage

Syntax

```
Set-AdSite -Identity ' SiteIdentity '
[-HubSiteEnabled <$true | $false>] [-InboundMailEnabled <$true | $false>]
[-DomainController ' DCName ' ] [-Name ' NewSiteName ' ]
```

Usage

```
Set-AdSite -Identity 'First-Seattle-Site' -HubSiteEnabled $true
```

Routing Messages Across Site Links

To view the configuration information about an Active Directory IP site link, you can use the `Get-AdSiteLink` cmdlet. This configuration information includes the value of the Exchange-specific cost, the cost assigned to the Active Directory IP site link, and a list of the sites in the IP site link.

Listing 20-3 provides the syntax and usage, as well as sample output, for the `Get-AdSiteLink` cmdlet. Use the `-Identity` parameter to retrieve the configuration information about a specific IP site link. If you do not provide an identity, the configuration information about all IP site links is returned.

LISTING 20-3 `Get-AdSiteLink` cmdlet syntax and usage

Syntax

```
Get-AdSiteLink [-Identity ' SiteIdentity ']  
[-DomainController ' DCName ']
```

Usage

```
Get-AdSiteLink -Identity 'PORTLANDSEATTLELINK' | fl
```

Output

```
RunspaceId      :  
Cost            : 100  
ADCost         : 100  
ExchangeCost   :  
MaxMessageSize : Unlimited  
Sites          : {imaginedlands.com/Configuration/Sites/  
First-Seattle-Site}  
AdminDisplayName :  
ExchangeVersion : 0.0 (6.5.6500.0)  
Name           : PORTLANDSEATTLELINK  
DistinguishedName : CN=PORTLANDSEATTLELINK,CN=IP,CN=Inter-Site  
                Transports,CN=Sites,CN=Configuration,  
                DC=pocket-consultant,DC=com  
Identity       : imaginedlands.com/Configuration/Sites/Inter-Site  
                Transports/IP/PORTLANDSEATTLELINK  
Guid          :  
ObjectCategory : imaginedlands.com/Configuration/Schema/Site-Link  
ObjectClass    : {top, siteLink}  
WhenChanged    : 2/15/2016 8:55:33 PM  
WhenCreated    : 2/15/2016 8:55:33 PM  
WhenChangedUTC : 2/16/2016 4:55:33 AM  
WhenCreatedUTC : 2/16/2016 4:55:33 AM  
OrganizationId :  
OriginatingServer : CorpServer27.imaginedlands.com  
IsValid        : True  
ObjectState    : Unchanged
```

By default, Exchange Server 2016 determines the least-cost routing path by using the cost that is assigned to the Active Directory IP site links. You can change this behavior by using the `Set-AdSiteLink` cmdlet to configure an Exchange-specific cost for Active Directory IP site links. After you configure it, the Exchange-specific cost is used to determine the Exchange routing path rather than the Active Directory–assigned cost.

Listing 20-4 provides the syntax and usage, for the `Set-AdSiteLink` cmdlet. When there are multiple wide area network (WAN) paths between sites, you can set a higher site-link cost to reduce the likelihood that a link will be used and a lower site-link cost to increase the likelihood that a link will be used. You must have Enterprise Administrator rights to use the `-Name` parameter to change the name of a site link.

You can use the `-MaxMessageSize` parameter to set the maximum size for messages that are relayed across a specified link. The default value is “unlimited,” which allows messages of any size to be relayed. You can specify the units for values by using B for bytes, KB for kilobytes, MB for megabytes, or GB for gigabytes. The valid range for maximum size is from 64 KB to the largest value in bytes that can be set using a 64-bit integer (9,223,372,036,854,775,807).

LISTING 20-4 Set-AdSiteLink cmdlet syntax and usage

Syntax

```
Set-AdSiteLink -Identity ' SiteIdentity '  
[-DomainController ' DCName '  
[-ExchangeCost Cost ]  
[-MaxMessageSize <' Size ' | 'Unlimited'>]  
[-Name ' NewSiteLinkName ']
```

Usage

```
Set-AdSiteLink -Identity 'PORTLANDSEATTLELINK'  
-ExchangeCost 20
```

```
Set-AdSiteLink -Identity 'LASACRAMENTOLINK'  
-MaxMessageSize 'Unlimited'
```

```
Set-AdSiteLink -Identity 'LASACRAMENTOLINK'  
-MaxMessageSize '256MB'
```

Managing Send Connectors

Send connectors are the gateways through which transport servers send messages, and only transport servers have Send connectors. Exchange automatically creates the Send connectors required for internal mail flow but does not create the Send connectors required for mail flow to the Internet. Send connectors are stored in Active Directory and are available to all transport servers in the Exchange organization by default.

Creating Send Connectors

As an administrator, you can explicitly create Send connectors for Internet mail flow and other necessary connectors, and then manage the configuration of these explicitly created Send connectors as needed. You cannot, however, manage the configuration of Send connectors created implicitly by Exchange to enable mail flow. The key reasons for creating Send connectors are to:

- Control explicitly how message routing works within domains or between domains.
- Control explicitly the hosts used as destinations or the way messages are routed over the Internet.
- Send mail to systems that are not Exchange servers.

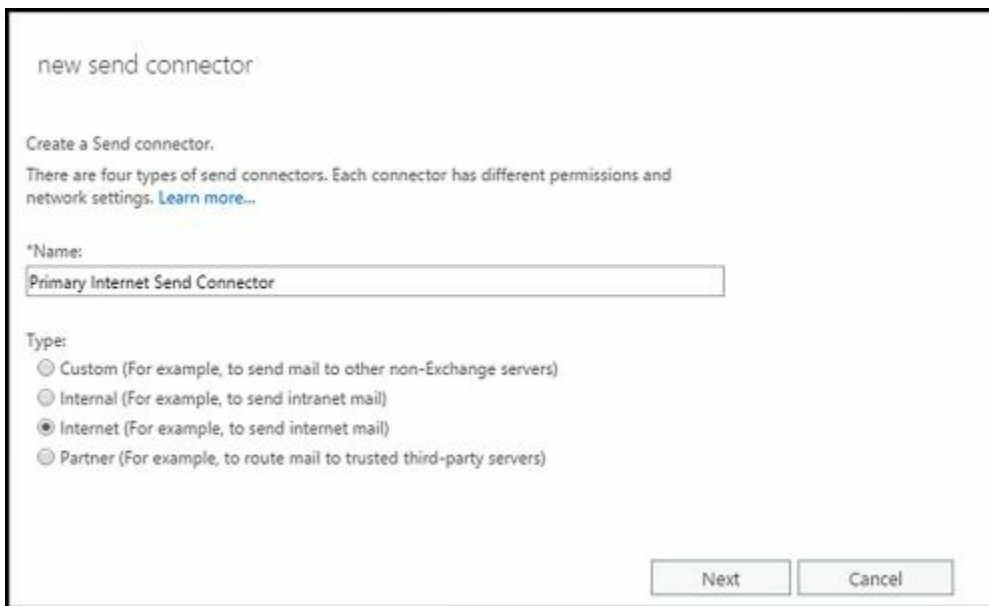
When you create Send connectors, you can encrypt message traffic sent over the link and require strict authentication. You can transmit messages to a designated internal server—called a *smart host*—or you can use DNS records to route messages. If you use a smart host, Exchange Server 2016 transfers messages directly to the smart host, which then sends out messages over an established link. The smart host allows you to route messages on a per-domain basis. If you use DNS records, Exchange Server 2016 performs a DNS lookup for each address to which the connector sends mail.

As part of the new architecture in Exchange 2016, Mailbox servers handle the front-end and back-end processing for mail transport, including the Transport service and the Front End Transport service. The Transport service is responsible for all mail flow, and the Front End Transport service acts as a stateless proxy for all external SMTP traffic. The Transport service on a Mailbox server can use a Send connector to route outbound messages through the Front End Transport service. Mail routing occurs internally. When an Active Directory site has subscribed Edge Transport servers, outbound mail is passed directly from a Mailbox server to an Edge Transport server.

When you create a Send connector, you must define the address space, which determines when the Send connector is used as well as the domain names to which the connector sends messages. For example, if you want to connect two domains in the same Exchange organization—`dev.tvpress.com` and `corp.tvpress.com`—you can create a Send connector in `dev.tvpress.com`, and then add an SMTP address type for the email domain `corp.tvpress.com`.

Send connectors can be used by multiple transport servers. When you create a Send

connector within an Exchange organization, you can specify the transport servers that are permitted to use the Send connector. When you create a Send connector on an Edge Transport server, the connector is configured only for that server.



new send connector

Create a Send connector.

There are four types of send connectors. Each connector has different permissions and network settings. [Learn more...](#)

*Name:
Primary Internet Send Connector

Type:

- Custom (For example, to send mail to other non-Exchange servers)
- Internal (For example, to send intranet mail)
- Internet (For example, to send internet mail)
- Partner (For example, to route mail to trusted third-party servers)

Next Cancel

FIGURE 20-1 Create a new SMTP Send connector for Internet mail flow.

Typically, the first Send connector you'll create in an Exchange organization is one that enables mail flow to the Internet. To create a Send connector for Internet mail flow, follow these steps:

1. In Exchange Admin Center, select Mail Flow in the Navigation menu and then select Send Connectors.
2. Click New. This starts the New Send Connector Wizard, shown in Figure 20-1.
3. In the Name text box, type a descriptive name for the connector, such as Primary Internet Send Connector, and then set the connector type as Internet. Click Next.
4. Confirm that MX Record Associated With Recipient Domain is selected, and then click Next.

new send connector

A send connector can route mail directly through DNS or redirect it to a smart host. [Learn more...](#)

*Network settings:
Specify how to send mail with this connector.

MX record associated with recipient domain
 Route mail through smart hosts

+ ✎ -

SMART HOST

Use the external DNS lookup settings on servers with transport roles

Back Next Cancel

5. You next need to define the address space for the send connector. Click the Add button (+).

new send connector

A Send connector routes mail to a specified list of domains. These domains can be an SMTP address space or a custom type. [Learn more...](#)

*Address space:
Specify the address space or spaces to which this connector will route mail.

+ ✎ -

TYPE	DOMAIN	COST

Scoped send connector

Back Next Cancel

6. In the Add Domain dialog box, SMTP is set as the address space type. In the Fully Qualified Domain Name box, enter * to specify that you are creating the connector for routing outbound mail to all external domains. By default, the address space cost is set to 1, which assigns the highest preference to the connector. If you plan to create other Send connectors, you may want to assign a higher cost to ensure mail is routed appropriately. For example, if you set the cost to 100 and the cost of other Send connectors to a value less than 100, this connector will be used only when no other connector would otherwise apply. Click Save to close the Add Domain dialog box. Click Next.


add domain

*Type:
SMTP

*Full Qualified Domain Name (FQDN):
*

*Cost:
1

Save Cancel

7. On the Source Server page, click the Add button () to associate the connector with the Mailbox server or servers that will be used to send mail to the Internet.

new send connector

A send connector sends mail from a list of servers with transport roles or Edge Subscriptions. [Learn more...](#)

*Source server:
Associate this connector with the following servers containing transport roles. You can also add Edge Subscriptions to this list.

+ -

SERVER	SITE	ROLE

Back Finish Cancel

8. In the Select A Server dialog box, select a Mailbox server that will be used as the source server, and then click Add. Repeat as necessary to add more Transport servers. If you make a mistake, click the Remove link next to the server name.

NAME	SITE	ROLE	VERSION
CORPSERVER74	imaginedlands.local/Configuratio...	Mailbox, ClientAccess	Version 15.1 (Build 2...

1 selected of 1 total

add -> CORPSERVER74[remove]:

OK Cancel

9. When you are finished selecting servers, click OK to close the Select A Server dialog box, and then click Finish to create the connector. You can verify that the

connector is configured properly by sending mail to an external recipient and confirming that the message arrives.

10. By default, the new Send connector is enabled and configured to allow a maximum message size of 35 MB. To change the default maximum message size, open the related properties dialog box by double-clicking the connector's entry in Exchange Admin Center. Next, enter the desired Maximum Send Message Size in the combo box provided and then click Save. Valid maximum send message sizes range from 1 to 2096128 MB. If you don't want the connector to have a specific limit, set the maximum size to 0.

To create other Send connectors, complete the following steps:

1. In Exchange Admin Center, select Mail Flow in the Navigation menu, and then select Send Connectors.
2. Click New to start the New Send Connector Wizard, shown in Figure 20-2.

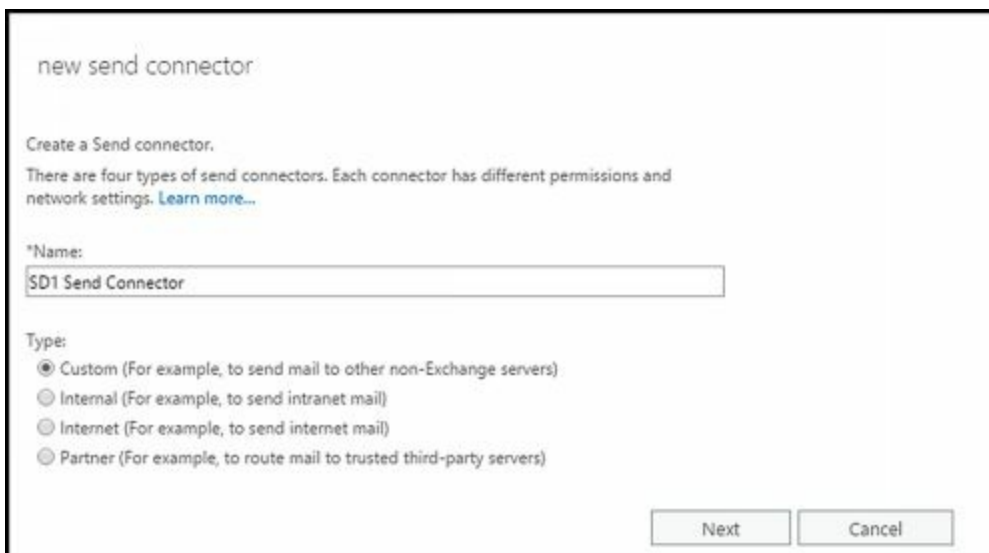
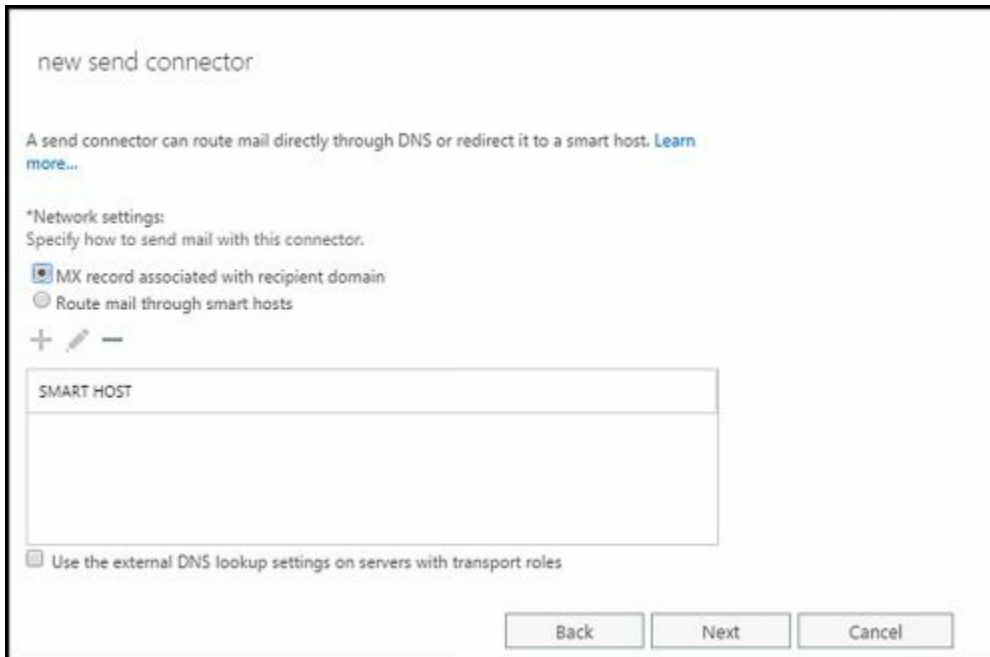


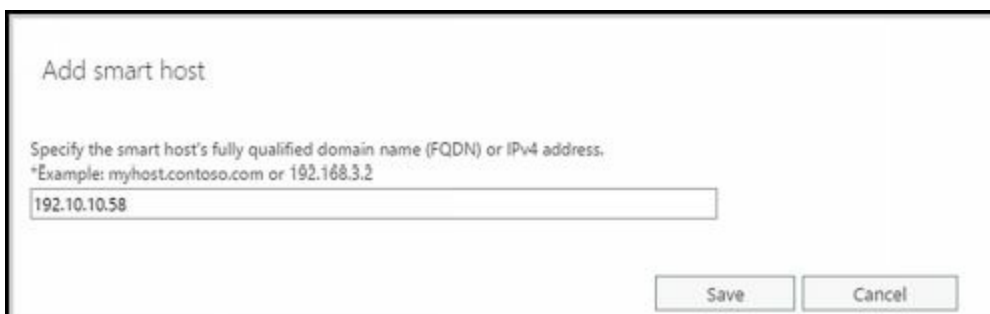
FIGURE 20-2 Create a new SMTP Send connector.

3. In the Name text box, type a descriptive name for the connector, and then set the connector type. The available options are as follows:
 - **Custom** Creates a customized Send connector for connecting with systems that are not Exchange servers.
 - **Internal** Creates a Send connector for sending mail to another transport server in the organization, and sets the default permissions so that the connector can be used by Exchange servers. This connector will be configured to route mail using smart hosts.
 - **Internet** Creates a Send connector that sends mail to external users over the Internet. This connector will be configured to use DNS records to route mail.
 - **Partner** Creates a Send connector that sends mail to partner domains. Partner domains cannot be configured as smart hosts. Only connections that authenticate with Transport Layer Security (TLS) certificates are allowed by default. Partner domains must also be listed on the TLS Send Domain Secure list, which can be set by using the `-TLSSendDomainSecureList` parameter of the `Set-TransportConfig` command.
4. On the Network Settings page, select how you want to send email with the Send

connector. If you select MX Record Associated With Recipient Domain, the Send connector uses the DNS client service on the Transport server to query a DNS server and resolve the destination address. Skip steps 5–9 if you select the MX Record option.



5. If you select Route Mail Through Smart Host, you have to specify the smart hosts to which mail should be forwarded for processing. Click the Add button (**+**).

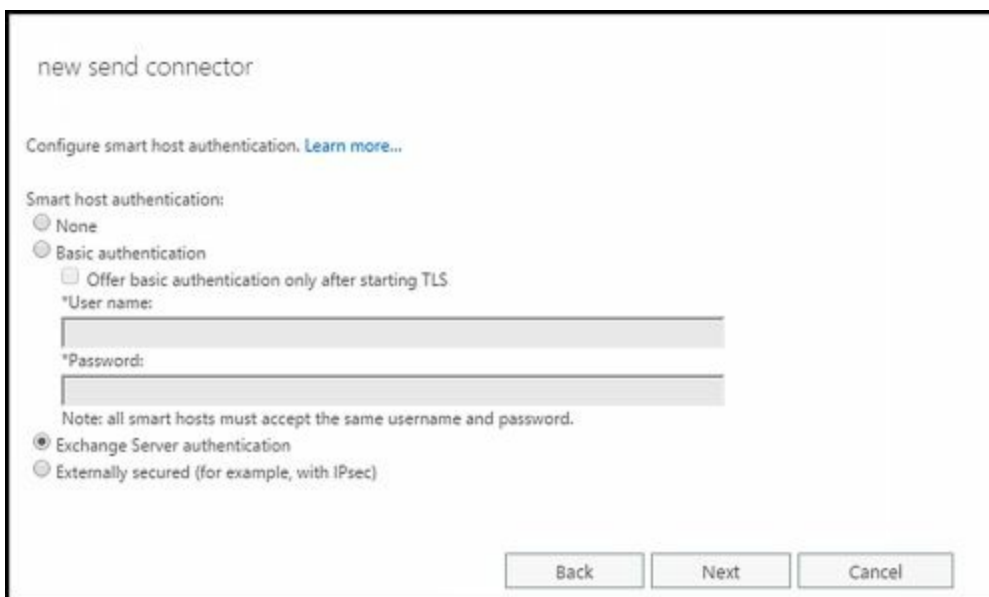


6. In the Add Smart Host dialog box, enter the IP address or the Fully Qualified Domain Name (FQDN) of the smart host. The Transport server must be able to resolve the FQDN.
7. Click Save to close the Add Smart Host dialog box. Repeat steps 5–7 as necessary to add more smart hosts to this connector. If you make a mistake, select the smart host, and then click Edit or Remove as appropriate. When you are finished, click Next to continue.
8. Optionally, specify that you want the connector to use the external DNS lookup settings of the Mailbox server.
9. After you've configured smart hosts, you'll see the Configure Smart Host Authentication Settings page. On this page, select the method that you want to use to authenticate your servers to the smart host. Choose one of the following options, and then click Next:
 - **None** No authentication. Use this option only if the smart host is configured to

accept anonymous connections.

- **Basic Authentication** Standard authentication with wide compatibility. With basic authentication, the user name and password specified are passed as cleartext to the remote domain.
- **Offer Basic Authentication Only After Starting TLS** When you use Basic Authentication, you can select this checkbox to enable basic authentication over TLS. In this case, TLS authentication is combined with basic authentication to allow encrypted authentication for servers with smart cards or X.509 certificates.
- **Exchange Server Authentication** Secure authentication for Exchange servers. With Exchange Server authentication, credentials are passed securely.
- **Externally Secured** Secure authentication for Exchange servers. With externally secured authentication, credentials are passed securely using an external security protocol for which the server has been separately configured, such as Internet Protocol security (IPSec).

NOTE With the Basic Authentication, you must provide the user name and password for the account authorized to establish connectors to the designated smart hosts. All smart hosts must use the same user name and password.



The screenshot shows a dialog box titled "new send connector" with the subtitle "Configure smart host authentication. Learn more...". Under "Smart host authentication:", there are three radio button options: "None", "Basic authentication", and "Exchange Server authentication". The "Basic authentication" option is selected, and it has a sub-option "Offer basic authentication only after starting TLS" which is unchecked. Below these options are two text input fields labeled "*User name:" and "*Password:". A note states "Note: all smart hosts must accept the same username and password." At the bottom, there are three buttons: "Back", "Next", and "Cancel".

10. Click Next. On the Address Space page, click the Add button ().

new send connector

A Send connector routes mail to a specified list of domains. These domains can be an SMTP address space or a custom type. [Learn more...](#)

*Address space:
Specify the address space or spaces to which this connector will route mail.

+ ✎ -

TYPE	DOMAIN	COST

Scoped send connector

Back Next Cancel

11. In the Add Domain dialog box, you can use the following options to specify the domain names to which this connector will send mail:

- **Type** SMTP is the default address space type. Use this type for connectors routing mail to Exchange server and other SMTP servers. For routing mail directly to non-SMTP servers, specify the address space type of the server, such as X400, X500, or MSMAIL.
- **Fully Qualified Domain Name** The domain or domains to which this connector will send mail, such as `imaginedlands.com`. With SMTP addresses, you can enter the wildcard character (*) directly in the address space as defined in RFC 1035. For example, you can enter * for all domains, *.com for all .com domains, or *.imaginedlands.com for the imaginedlands.com domain and all subdomains of imaginedlands.com. With X.400 addresses, you must specify the address space as defined in RFC 1685.
- **Cost** The address space cost is used for relative weighting. Valid address space costs range from 1, which assigns the highest possible preference, to 100, which assigns the lowest possible preference. When you create a Send connector, the default address space cost is 1. If you set all address spaces to this cost, all address spaces have equal preference for routing mail.

add domain

*Type:
SMTP


*Full Qualified Domain Name (FQDN):
*.imaginedlands.com

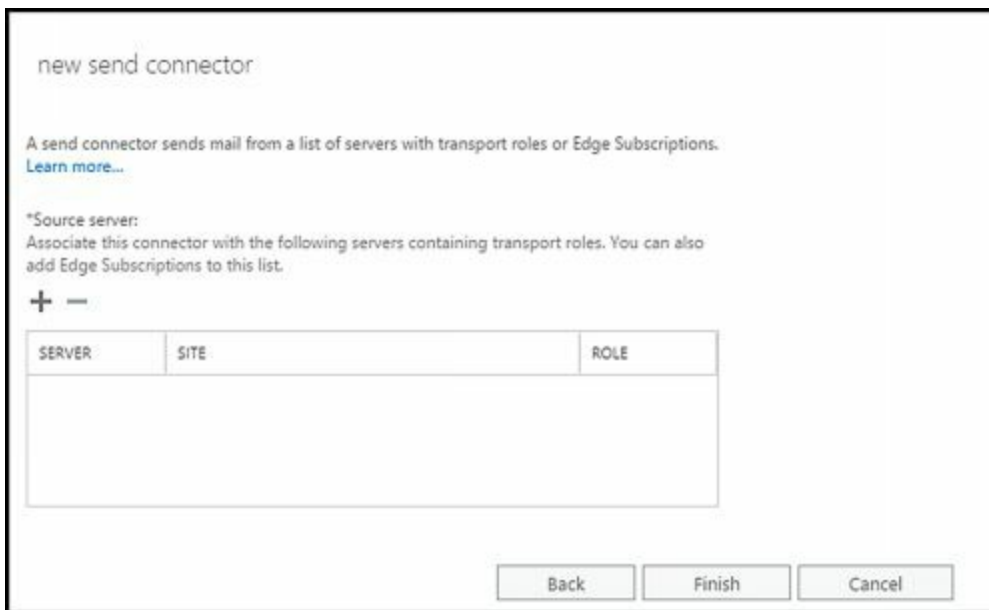
*Cost:
1

Save Cancel

12. Click Save to close the SMTP Address Space dialog box. Repeat as necessary to add more address spaces to this connector. If you make a mistake, select the

address space and then click Edit or Remove as appropriate.

13. If you'd like to scope the Send connector to the current site, select the Scoped Send Connector check box. When a Send connector is scoped, only Mailbox servers in the same Active Directory site as the Send connector's source servers consider that Send connector in routing decisions. Click Next to continue.
14. Next, you see the Source Server page, allowing you to associate this connector with other transport servers. Click the Add button () to associate the connector with Mailbox servers and Edge subscriptions.



new send connector

A send connector sends mail from a list of servers with transport roles or Edge Subscriptions. [Learn more...](#)

*Source server:
Associate this connector with the following servers containing transport roles. You can also add Edge Subscriptions to this list.

+ -

SERVER	SITE	ROLE
--------	------	------

Back Finish Cancel

15. In the Select A Server dialog box, select a Mailbox server or an Edge subscription that will be used as the source server for sending messages to the address space that you previously specified, and then click Add. Repeat as necessary to add more Transport servers. If you make a mistake, click the Remove link next to the server name.



NAME	SITE	ROLE	VERSION
CORPSERVER74	imaginedlands.local/Configuratio...	Mailbox, ClientAccess	Version 15.1 (Build 2...

1 selected of 1 total

add -> CORPSERVER74[remove:]

OK Cancel

16. When you are finished, click OK to close the Select A Server dialog box, and then click Finish to create the connector. You can verify that the connector is configured properly by sending mail to an external recipient in a domain

associated with the connector and confirming that the message arrives.

17. By default, the new Send connector is enabled and configured to allow a maximum message size of 35 MB. To change the default maximum message size, open the related properties dialog box by double-clicking the connector's entry in Exchange Admin Center. Next, enter the desired Maximum Send Message Size in the combo box provided, and then click Save. Valid maximum send message sizes range from 1 to 2048 MB. If you don't want the connector to have a specific limit, set the maximum size to 0 or select Unlimited in the dropdown list.
18. When you create a Send connector, you can also enable front-end proxying to allow destination specific routing, such as by DNS or IP address. If you want to enable front-end proxying and route outbound messages through the Client Access services running on a Mailbox server in the local Active Directory site, open the related properties dialog box by double-clicking the connector's entry in Exchange Admin Center. Next, select the Proxy Through Client Access Services check box, and then click Save.

In Exchange Management Shell, you can create Send connectors by using the `New-SendConnector` cmdlet. The `-Usage` parameter sets the Send connector type as Custom, Internal, Internet, or Legacy. The `-AddressSpaces` parameter sets the address spaces for the Send connector by FQDN or IP address. Additionally, the `-DNSRoutingEnabled` parameter determines whether DNS records or smart hosts are used for lookups. To use DNS records, set `DNSRoutingEnabled` to `$true`. To use smart hosts, set `-DNSRoutingEnabled` to `$false`, and then use the `-SmartHosts` parameter to designate the smart hosts. To enable front-end proxying, set the `-FrontEndProxyEnabled` parameter to `$true`.

Listing 20-5 provides the syntax and usage for the `New-SendConnector` cmdlet. With basic authentication or basic authentication over TLS, you will be prompted to provide credentials. To scope the Send connector to the current Active Directory site, set the `-IsScopedConnector` parameter to `$true`.

LISTING 20-5 `New-SendConnector` cmdlet syntax and usage

Syntax

```
New-SendConnector -Name Name -AddressSpaces Addresses
[-AuthenticationCredential Credentials ]
[-CloudServicesMailEnabled <$true | $false>]
[-Comment Comment ]
[-ConnectionInactivityTimeout TimeSpan ]
[-Custom <$true | $false>]
[-DNSRoutingEnabled <$true | $false>]
[-DomainController DCName ]
[-DomainSecureEnabled <$true | $false>]
[-ErrorPolicies <Default|DowngradeDnsFailures|DowngradeCustomFailures>]
[-Enabled <$true | $false>]
[-Force <$true | $false>]
[-ForceHELO <$true | $false>]
```

```
[-Fqdn FQDN ]
[-FrontEndProxyEnabled <$true | $false>]
[-IgnoreStartTLS <$true | $false>]
[-Internal <$true | $false>]
[-Internet <$true | $false>]
[-IsScopedConnector <$true | $false>]
[-MaxMessageSize < Size | Unlimited>]
[-Partner <$true | $false>]
[-Port PortNumber ]
[-ProtocolLoggingLevel <None | Verbose>]
[-RequireTLS <$true | $false>]
[-SmartHostAuthMechanism <None|BasicAuth|BasicAuthRequireTLS
    |ExchangeServer|ExternalAuthoritative>]
[-SmartHosts SmartHosts ]
[SmtptMaxMessagesPerConnection MaxMessages ]
[-SourceIPAddress IPAddress ]
    [-SourceTransportServers TranportServers ]
[-TlsAuthLevel <EncryptionOnly|CertificateValidation|DomainValidation>]
[-TlsCertificateName "X509:<I> Issuer <S> CommonName "]
[-TlsDomain DomainNameForVerificationofTLSCert ]
[-Usage <Custom|Internal|Internet|Partner>]
[-UseExternalDNSServersEnabled <$true | $false>]
```

Usage for DNS MX records

```
New-SendConnector -Name "Imaginedlands.com Send Connector"
-Usage "Custom"
-AddressSpaces "smtp:*.imaginedlands.com;1"
-IsScopedConnector $true
-DNSRoutingEnabled $true
-UseExternalDNSServersEnabled $false
-SourceTransportServers "CORPSVR127"
```

Usage for smart hosts

```
New-SendConnector -Name "Cohovineyards.com"
-Usage "Custom"
-AddressSpaces "smtp:*.cohovineyards.com;1"
-IsScopedConnector $false
-DNSRoutingEnabled $false
-SmartHosts "[192.168.10.52]"
-SmartHostAuthMechanism "ExternalAuthoritative"
-UseExternalDNSServersEnabled $false
-SourceTransportServers "CORPSVR127"
```

Viewing and Managing Send Connectors

The Exchange Management tools provide access only to the Send connectors you've explicitly created. On Mailbox servers, Send connectors created by Exchange Server are not displayed or configurable. On Edge Transport servers, you can view and manage the internal Send connector used to connect to the Mailbox servers in your Exchange

organization, as shown in Figure 20-3.

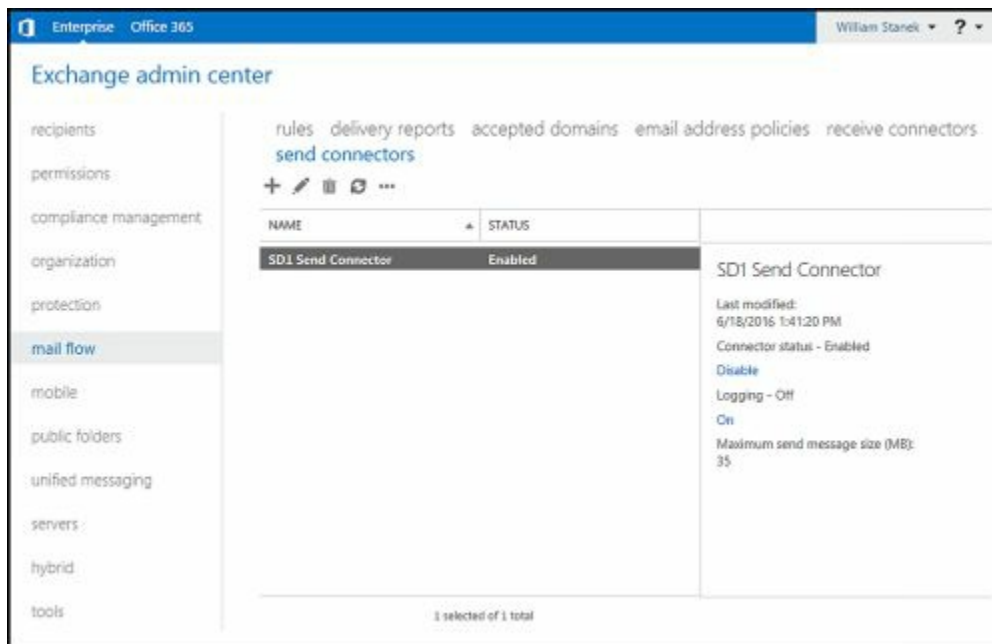



FIGURE 20-3 Viewing Send connectors in your on-premises Exchange organization.

In Exchange Admin Center, you can view the Send connectors and manage their configuration. When you select Mail Flow in the Navigation menu and then select Send Connectors, Send connectors you've created are listed by name and status. You can now do the following:

- **Change a connector's properties** To change a connector's properties, double-click the connector entry, and then use the Properties dialog box to manage the connector's properties. You'll also be able to specify the maximum message size and protocol logging level. By default, the maximum message size is set to 35 MB and the protocol logging level is set to None.
- **Enable a connector** To enable a connector, select it, and then select Enable in the details pane.
- **Disable a connector** To disable a connector, select it, and then select Disable in the details pane.
- **Remove a connector** To remove a connector, select it in Exchange Admin Center, and then select Delete ().

In Exchange Management Shell, you can view, update, or remove Send connectors by using the Get-SendConnector, Set-SendConnector, or Remove-SendConnector cmdlets, respectively. Listings 20-6 through 20-8 provide the syntax and usage. With Get-SendConnector, if you don't specify an identity, the cmdlet returns a list of all administrator-configured Send connectors.

LISTING 20-6 Get-SendConnector cmdlet syntax and usage

Syntax

Get-SendConnector

Get-SendConnector –Identity **ConnectorIdentity**

[-DomainController **DCName**]

Usage

Get-SendConnector -Identity "Imaginedlands.com Send Connector"

LISTING 20-7 Set-SendConnector cmdlet syntax and usage

Syntax

Set-SendConnector –Identity **ConnectorIdentity**

[-Name **NewName**]

[-AddressSpaces **Addresses**]

[-AuthenticationCredential **Credentials**]

[-CloudServicesMailEnabled <\$true | \$false>]

[-Comment **Comment**]

[-ConnectionInactivityTimeout **TimeSpan**]

[-DNSRoutingEnabled <\$true | \$false>]

[-DomainController **DCName**]

[-DomainSecureEnabled <\$true | \$false>]

[-ErrorPolicies <Default|DowngradeDnsFailures|DowngradeCustomFailures>]

[-Enabled <\$true | \$false>]

[-Force <\$true | \$false>]

[-ForceHELO <\$true | \$false>]

[-Fqdn **FQDN**]

[-FrontEndProxyEnabled <\$true | \$false>]

[-IgnoreStartTLS <\$true | \$false>]

[-IsScopedConnector <\$true | \$false>]

[-MaxMessageSize < **Size** | Unlimited>]

[-Port **PortNumber**]

[-ProtocolLoggingLevel <None | Verbose>]

[-RequireTLS <\$true | \$false>]

[-SmartHostAuthMechanism <None|BasicAuth|BasicAuthRequireTls
|ExchangeServer|ExternalAuthoritative>]

[-SmartHosts **SmartHosts**]

[-SourceIPAddress **IPAddress**]

[-SourceTransportServers **TransportServers**]

[SmtpMaxMessagesPerConnection **MaxMessages**]

[-TlsAuthLevel <EncryptionOnly|CertificateValidation|DomainValidation>]

[-TlsCertificateName "X509:<I> **Issuer** <S> **CommonName** "]

[-TlsDomain **DomainNameForVerificationofTLSCert**]

[-UseExternalDNSServersEnabled <\$true | \$false>]

Usage

Set-SendConnector -Identity "Imaginedlands.com Send Connector"

-AddressSpaces "smtp:*.imaginedlands.com;1"

-DNSRoutingEnabled \$true -SmartHosts 10.10.2.205

-SmartHostAuthMechanism "None"

-SourceTransportServers "CORPSVR127"

LISTING 20-8 Remove-SendConnector cmdlet syntax and usage

Syntax

```
Remove-SendConnector -Identity ConnectorIdentity  
[-Confirm <$true | $false>] [-DomainController DCName ]
```

Usage

```
Remove-SendConnector -Identity "Imaginedlands.com Send Connector"
```

Configuring Send Connector DNS Lookups

You can configure different settings for internal and external DNS lookups by configuring a Transport server's External DNS Lookups and Internal DNS Lookups properties. External DNS Lookup servers are used to resolve the IP addresses of servers outside your organization. Internal DNS Lookup servers are used to resolve IP addresses of servers inside the organization.

In Exchange Admin Center, you can specify enable or disable external DNS lookups for each Send connector by selecting Mail Flow in the Navigation menu, and then selecting Send Connectors. Next, double-click the Send connector you want to configure. In the properties dialog box, select Delivery to display the Delivery options. The Use The External DNS Lookup... checkbox controls whether external DNS lookups are permitted. To allow external DNS lookups when the selected connector is used, select this checkbox, and then click Save.

If you've enabled external DNS lookups for Send connectors, you can specify how external lookups should be performed for each Mailbox server in the organization. You also can configure internal DNS lookups for each Mailbox server in the organization. To configure DNS Lookup servers, complete these steps:

1. In Exchange Admin Center, select Servers in the Navigation menu, and then select Servers. Next, double-click the server you want to manage.
2. In the properties dialog box, select DNS Lookups to display DNS lookup options.
3. On the External DNS Lookups panel, shown in Figure 20-4, specify how external lookups should be performed:
 - To use DNS settings from the server's network card or cards for external lookups, choose either All Network Adapters (All Available IPv4) to use all configured IPv4 settings or a specific network card to use the configured IPv4 settings of that card.
 - To use a custom list of DNS servers for external lookups, select Custom Settings. Next, click Add. In the Add IP Address dialog box, type the IPv4 or IPv6 address of a DNS server to use for external lookups, and then click Save. Repeat this process to specify multiple servers. Keep in mind that Mailbox servers perform lookups in the order the DNS servers are listed.

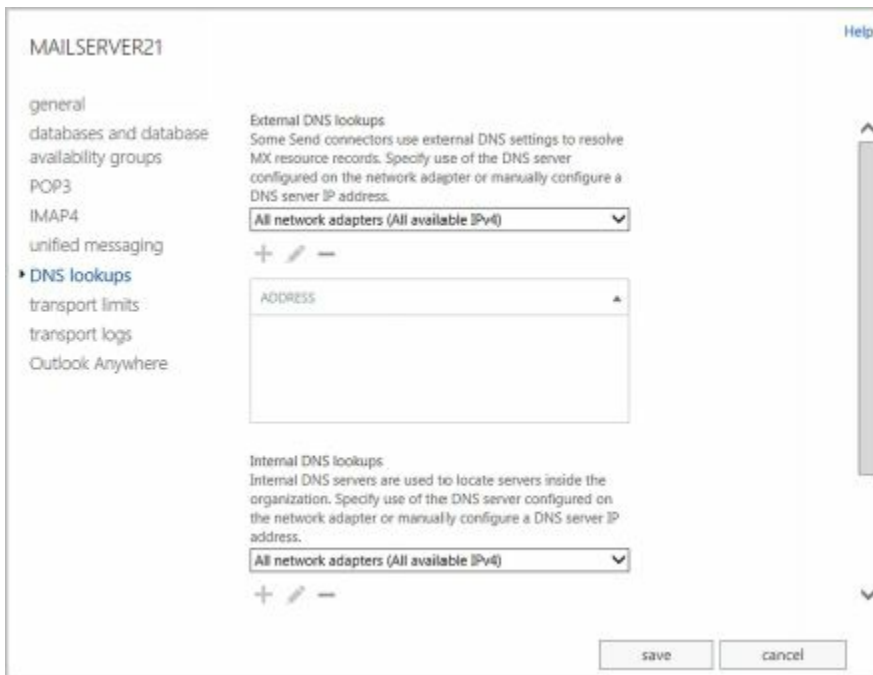


FIGURE 20-4 Configure external DNS lookups.

4. On the Internal DNS Lookups panel, specify how internal lookups should be performed:
 - To use DNS settings from the server’s network card or cards for internal lookups, choose either All Network Adapters (All Available IPv4) to use all configured IPv4 settings or a specific network card to use the configured IPv4 settings of that card.
 - To use a custom list of DNS servers for internal lookups, select Custom Settings. Next, click Add. In the Add IP Address dialog box, type the IPv4 or IPv6 address of a DNS server to use for internal lookups, and then click Save. Repeat this process to specify multiple servers.
5. Click Save to apply your settings.

Setting Send Connector Limits

Send connector limits determine how mail is delivered after a connection has been established and the receiving computer has acknowledged that it’s ready to receive the data transfer. After a connection has been established and the receiving computer has acknowledged that it’s ready to receive the data transfer, Exchange Server attempts to deliver messages queued for delivery to the computer. If a message can’t be delivered on the first attempt, Exchange Server tries to send the message again after a specified time. Exchange Server keeps trying to send the message at the intervals you’ve specified until the expiration time-out is reached. When the time limit is reached, the message is returned to the sender with a non-delivery report (NDR). The default expiration time-out is two days.

After multiple failed attempts to deliver a message, Exchange Server generates a delay notification and queues it for delivery to the sender of the message. Notification doesn’t occur immediately after failure; instead, Exchange Server sends the delay notification message after the notification delay interval and then only if the message hasn’t already

been delivered. The default delay notification is four hours.

With SMTP, you have much more control over outgoing connections than you do over incoming connections. You can limit the number of simultaneous connections and the number of connections per domain. These limits set the maximum number of simultaneous outbound connections. By default, the maximum number of connections is 1,000 and the maximum number of connections per domain is 20.

You can view or change the Send connector limits by completing the following steps:

1. In Exchange Admin Center, select Servers in the Navigation menu, and then select Servers. Next, double-click the server you want to manage.
2. On the Transport Limits page, shown in Figure 20-5, use the following options for retrying unsuccessful outbound connections:
 - **Outbound Connection Failure Retry Interval (Seconds)** Sets the retry interval for subsequent connection attempts to a remote server where previous connections have failed. The default is 600 seconds.
 - **Transient Failure Retry Interval (Minutes)** Sets the interval at which the server immediately retries when it encounters a connection failure with a remote server. The default is five minutes.
 - **Transient Failure Retry Attempts** Sets the maximum number of times that the server immediately retries when it encounters a connection failure with a remote server. The default is six. If you enter 0 as the number of retry attempts or the maximum number of attempts has been reached, the server no longer immediately retries a connection and instead waits according to the outbound connection failure retry interval.
3. When messages that cannot be delivered reach the Maximum Time Since Submission value, they expire, and Exchange Server generates a Non-delivery report. To set the expiration time-out for messages, enter the desired message expiration value in the Maximum Time Since Submission (Days) text box. The default expiration time-out for messages is two days.

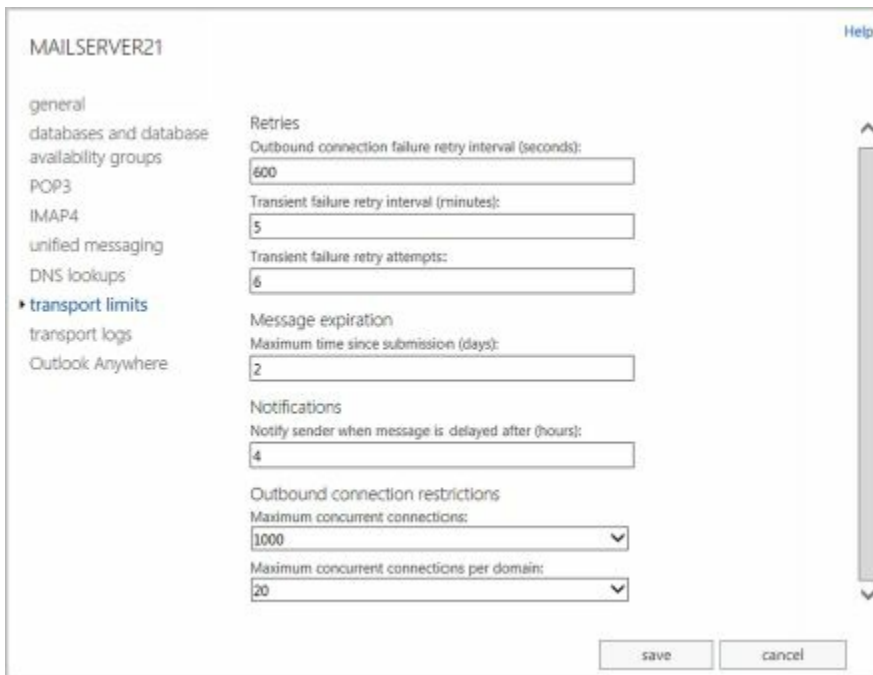


FIGURE 20-5 Configure connection limits.

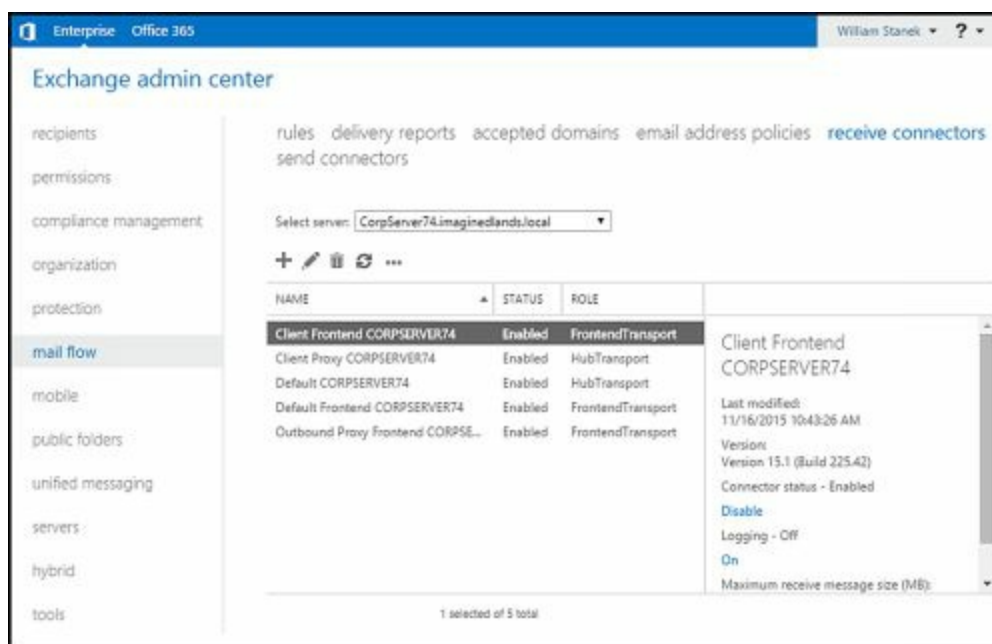
4. When messages are delayed longer than the allowed delay interval, Exchange Server sends a delay notification to the sender. To set the amount of time to wait before notifying senders of a delay, enter the desired wait time in the Notify Sender When Message Is Delayed After (Hours) text box. The default wait time is four hours.
5. To set an outgoing connection limit, select the Maximum Concurrent Outbound Connections check box, and then type the limit value. The default limit is 1,000 outbound connections. To remove outgoing connection limits, set the value to 0 or select Unlimited in the drop-down list.
6. To set an outgoing connection limit per domain, select the Maximum Concurrent Outbound Connections Per Domain check box, and then type the limit value. The default limit is 20 outbound connections per domain. To remove the outgoing connection limit per domain, set the value to 0 or select Unlimited in the drop-down list.
7. Click Save to apply your settings.

Managing Receive Connectors

Receive connectors are the gateways through which transport servers receive messages. Exchange creates the Receive connectors required for mail flow automatically. The receive permissions on a Receive connector determine who is allowed to send mail through the connector.

Two Receive connectors are created for front-end transport and two for hub transport on each Exchange 2016 Mailbox server:

- The Default connector for hub transport accepts connections from Mailbox servers running the Transport service as well as from Edge Transport servers
- The Client Proxy connector for hub transport accepts connections from Client Access services.
- The Default front-end connector accepts connections from SMTP senders over port 25.
- The Client front-end connector accepts secure connections over TLS
- The Outbound Proxy front-end connector accepts connections from Mailbox servers when front-end proxying is enabled.



As these default Receive connectors are created automatically, you generally don't need to create Receive connectors to receive mail from the Internet. That said, as an administrator, you can explicitly create Receive connectors, and then manage the configuration of those Receive connectors as necessary. You cannot, however, manage the configuration of connectors created implicitly by Exchange to enable mail flow. The key reasons for creating SMTP connectors are when you want to:

- Control explicitly how messages are received within domains or between domains.
- Control explicitly the permitted incoming connections.
- Receive mail from systems that are not Exchange servers.

Unlike Send connectors, Receive connectors are used by only a single, designated Transport server. When you create a Receive connector within an Exchange organization, you can select the Mailbox or Edge Transport server with which the connector should be associated and configure the specific binding for that connector. A binding is a combination of local IP addresses, ports, and remote IP address ranges for the Receive connector. You cannot create a Receive connector that duplicates the bindings of existing Receive connectors. Each Receive connector must have a unique binding.

NOTE Exchange Server 2016 uses standard SMTP or Extended SMTP (ESMTP) to deliver mail. Because the ESMTP standard is more efficient and allows for extensions, SMTP connectors always try to initiate ESMTP sessions before trying to initiate standard SMTP sessions. SMTP connectors initiate ESMTP sessions with other mail servers by issuing an EHLO start command. SMTP connectors initiate SMTP sessions with other mail servers by issuing the HELO start command.

SMTP was originally defined in RFC 821, and ESMTP was originally defined in RFC 1869. With SMTP, the MAIL FROM and RCPT TO fields are limited to a maximum of 512 characters. With ESMTP, these fields can have more than 512 characters. Additionally, EHLO replies can include a status code, domain, and a list of keywords that indicate supported extensions.

Because the ESMTP standard is more efficient and allows for extensions SMTP connectors always try to initiate ESMTP sessions before trying to initiate standard SMTP sessions. SMTP connectors initiate ESMTP sessions with other mail servers by issuing an EHLO start command. SMTP connectors initiate SMTP sessions with other mail servers by issuing the HELO start command.


Creating Receive Connectors

Receive connectors can be configured for either front-end transport or hub transport. Generally, when you want to control mail flow from external sources, you configure the Receive connector for front-end transport rather than hub transport. Thus, you normally would:

- [Configure a Receive connector for receiving messages from the Internet or an external partner as a front end transport.](#)
- [Configure a Receive connector for receiving messages from an internal messaging appliance or an internal Exchange server as a hub transport.](#)

You can create a Receive connector for front-end or hub transport by completing the following steps:

1. In Exchange Admin Center, select Mail Flow in the Navigation menu, and then select Receive Connectors. On the Receive Connectors page, use the Select Server list to choose the server on which you want to create the Receive connector.

2. Click New () to start the New Receive Connector Wizard, shown in Figure 20-6. In the Name text box, type a descriptive name for the connector, and then specify the connector role as either Hub Transport or Frontend Transport.

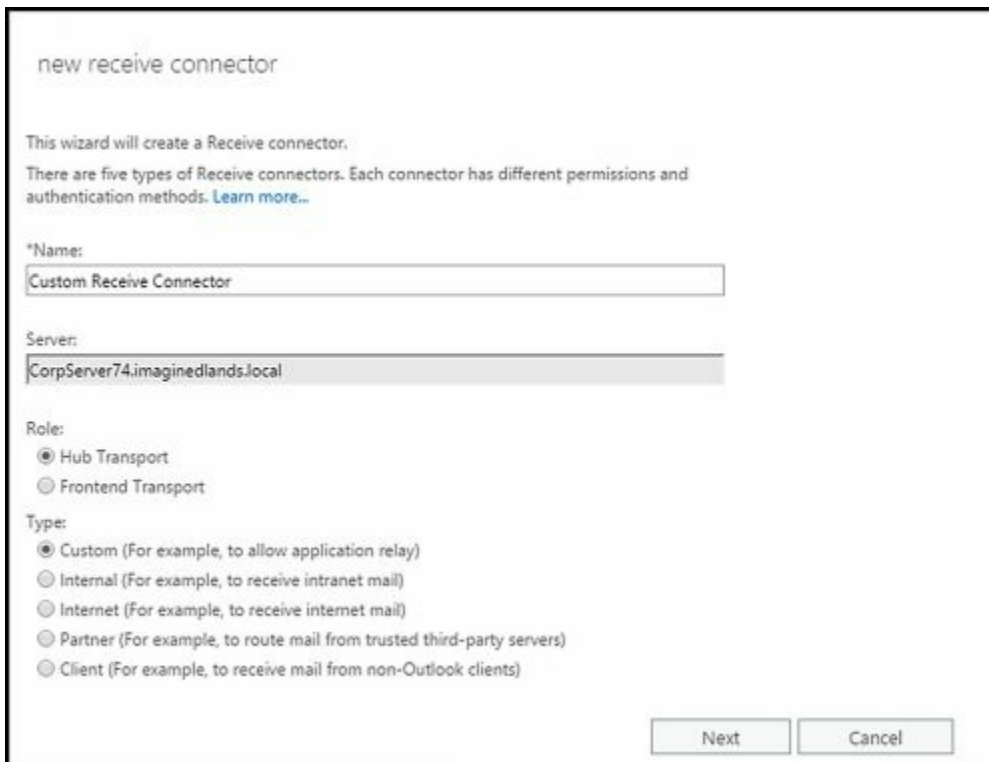

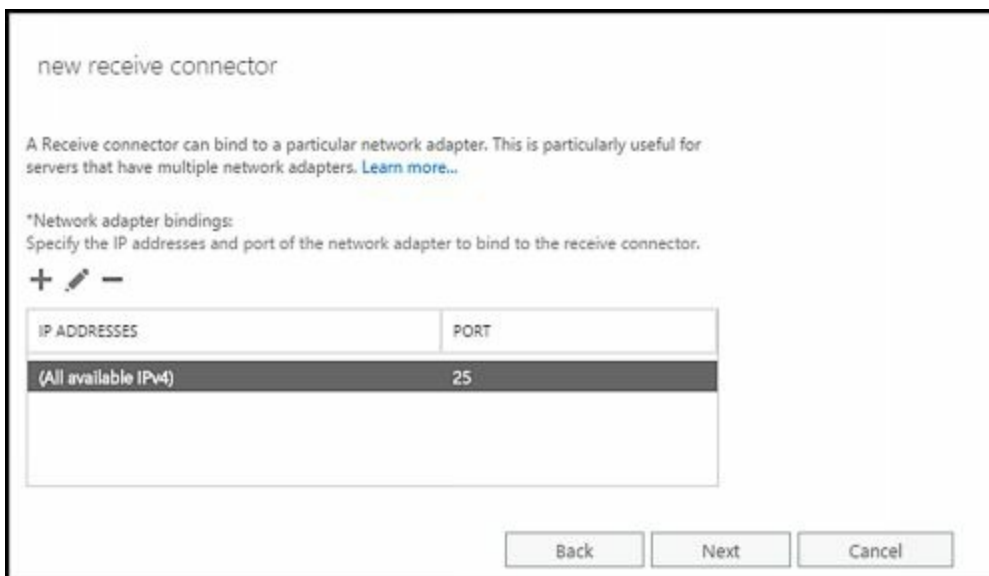



FIGURE 20-6 Create a new SMTP Receive connector.


3. Set the connector type. The available options are as follows:
 - **Custom** Creates a Receive connector bound to a specific port or IP address on a server with multiple receive ports or IP addresses. It can also be used to specify a remote IP address from which the connector receives messages. Generally, a custom Receive connector is used to connect with systems that are not Exchange servers. You also can use custom Receive connectors to receive mail from a Mailbox server in another forest or from an SMTP transfer agent.
 - **Internal** Creates a Receive connector to receive messages from another Transport server in the organization, such as may be necessary for communication between Mailbox servers or between Mailbox servers and third-party transfer agents. For Edge Transport servers, the Internal connector type sets the default permissions so that the connector can be used by Exchange servers. For Mailbox servers, it sets the default permissions so that the connector is configured to accept connections from Exchange servers.
 - **Internet** Creates a Receive connector that accepts incoming connections from the Internet. This connector accepts connections from anonymous users.
 - **Client** Creates a Receive connector used to receive mail from Exchange users. Only connections from authenticated Microsoft Exchange users are accepted by default. Typically used to connect clients not using Microsoft Office Outlook.
 - **Partner** Creates a Receive connector used to receive mail from partner domains. Partner domains cannot be configured as smart hosts. Only connections that


authenticate with Transport Layer Security (TLS) are allowed by default. Partner domains must also be listed on the TLS Receive Domain Secure list, which can be set by using the `-TLSReceiveDomainSecureList` parameter of the `Set-TransportConfig` command.

4. Click Next. For Custom, Partner, and Internet Receive connectors, you can specify the local IPv4 and IPv6 addresses and the port on which mail can be received. By default, Custom and Internet Receive connectors are configured to receive mail over port 25 on all available IPv4 addresses configured for the server. Port 25 is the default TCP port for SMTP. To use a different configuration for IPv4 addresses, select the default entry on the Local Network Settings page, and then click Remove ().



5. To create a new entry, click Add (). In the Add IP Address dialog box, select All Available IPv4 Addresses to have the connector listen for connections on all the IPv4 addresses that are assigned to the network adapters on the local server. Alternatively, you can select all available IPv6 addresses or you can select Specify An IPv4 Address Or an IPv6 Address if you want to type an IP address that is assigned to a network adapter on the local server and have the connector listen for connections only on this IP address. As necessary, modify the listen port value. Click Save. Repeat this process as necessary. When you are ready to continue, click Next.

- On the Remote Network Settings page, you can specify the remote IP addresses from which the server can receive mail. By default, Receive connectors are configured to accept mail from all remote IP addresses, which is why the IP address range 0.0.0.0–255.255.255.255 is set as the default entry. You'll only want to change this behavior if you want to limit the servers that are permitted to send mail to the Transport server. To use a different configuration, select the default entry on the Remote Network Settings page, and then click Remove ().

- To specify the remote servers, click Add (). Next, in the Add IP Address dialog box, enter an IP address, an IP address range, or an IP address range in Classless Internet Domain Routing (CIDR) notation. Repeat this process as necessary to specify other acceptable IP addresses. Click Save.
- When you're finished, click Finish to create the connector. You can verify that the connector is configured properly by confirming that messages arrive from a sending server to which the connector applies.
- By default, the new Receive connector is enabled and configured to allow a maximum message size of 35 MB. To change the default maximum message size, open the related properties dialog box by double-clicking the connector's entry in Exchange Admin Center. Next, enter the desired Maximum Receive Message Size

in the combo box provided, and then click Save. Valid maximum receive message sizes range from 1 to 2047 MB--and you can't specify that there is no limit.

10. When you create a Receive connector, you can also specify the maximum hops that a message can take before it's rejected by the Receive connector. By default, a message can have a maximum of 12 local hops and a maximum of 60 hops in total. If you want to change the default maximum hop counts, open the related properties dialog box by double-clicking the connector's entry in Exchange Admin Center. After you set the maximum number of local hops and the maximum number of hops in total, click Save. The valid range for local hops is 1 to 50 and the valid range for hops in total is 1 to 500. If you don't want the connector to have a specific limit for local hops, set the maximum local hops to 0. You can't set the maximum hops in total to unlimited.

In Exchange Management Shell, you can create Receive connectors by using the `New-ReceiveConnector` cmdlet. The `-Usage` parameter sets the Receive connector type as Client, Custom, Internal, Internet, or Partner. The `-Bindings` parameter sets the internal IP addresses and ports on which to listen. The `-FQDN` parameter sets the FQDN to advertise in response to HELO or EHLO messages. The `-RemoteIPRanges` parameter provides a comma-separated list of acceptable IP address ranges. To specify the server on which to create the Receive connector, use the `-Server` parameter.

As Listing 20-9 shows, the required parameters for the `New-ReceiveConnector` cmdlet depend on the type of Receive connector you are creating. After you provide the required parameters, the remaining parameters can be used in the same way regardless of which type of Receive connector you are creating. To specify the authentication type, use `-AuthMechanism`. With Basic Authentication or Basic Authentication Over TLS, you will be prompted to provide credentials. Also, you can use `-TransportRole` to specify the role with which the Receive connector should be associated.

LISTING 20-9 `New-ReceiveConnector` cmdlet syntax and usage

Syntax

```
New-ReceiveConnector -Name Name  
-Usage <Custom | Internet | Internal | Client | Partner> {AddtlParams}  
[-TransportRole <FrontEndTransport | HubTransport>]
```

```
New-ReceiveConnector -Name Name -Bindings Bindings  
-RemoteIPRanges IPRange1 , IPRange2 , . . . {AddtlParams}  
[-TransportRole <FrontEndTransport | HubTransport>]
```

```
New-ReceiveConnector -Name Name -Bindings Bindings  
-Internet <$true | $false > {AddtlParams}  
[-TransportRole <FrontEndTransport | HubTransport>]
```

```
New-ReceiveConnector -Name Name -Client <$true | $false >  
-RemoteIPRanges IPRange1 , IPRange2 , . . . {AddtlParams}  
[-TransportRole <FrontEndTransport | HubTransport>]
```


New-ReceiveConnector -Name **Name** -Internal <\$true | \$false >
-RemoteIPRanges **IPRange1** , **IPRange2** , . . . {AddtlParams}
[-TransportRole <FrontEndTransport | HubTransport>]

New-ReceiveConnector -Name <String> -Bindings **Bindings**
-Partner <\$true | \$false > -RemoteIPRanges **IPRange1** , **IPRange2** , . . .
[-TransportRole <FrontEndTransport | HubTransport>]
{AddtlParams}

{AddtlParams}
[-AdvertiseClientSettings <\$true | \$false>]
[-AuthMechanism <None | Tls | Integrated | BasicAuth |
BasicAuthRequireTLS | ExchangeServer | ExternalAuthoritative>]
[-Banner **Banner**]
[-BinaryMimeEnabled <\$true | \$false>]
[-Bindings **Bindings**]
[-ChunkingEnabled <\$true | \$false >]
[-Comment **Comment**]
[-ConnectionInactivityTimeout **TimeSpan**]
[-ConnectionTimeout **TimeSpan**]
[-Custom <\$true | \$false >]
[-DefaultDomain **DefaultDomain**]
[-DeliveryStatusNotificationEnabled <\$true | \$false>]
[-DomainController **DCName**]
[-DomainSecureEnabled <\$true | \$false>]
[-EightBitMimeEnabled <\$true | \$false>]
[-EnableAuthGSSAPI <\$true | \$false>]
[-Enabled <\$true | \$false>]
[-EnhancedStatusCodesEnabled <\$true | \$false>]
[-ExtendedProtectionPolicy <none | allow | require>]
[-Fqdn **FQDN**]
[-LiveCredentialEnabled <\$true | \$false>]
[-LongAddressesEnabled <\$true | \$false>]
[-MaxAcknowledgementDelay **MaxDelay**]
[-MaxHeaderSize **MaxHeaderBytes**]
[-MaxHopCount **MaxHops**]
[-MaxInboundConnection < **MaxConn** | Unlimited>]
[-MaxInboundConnectionPercentagePerSource **MaxPercentage**]
[-MaxInboundConnectionPerSource < **MaxConnPerSource** | Unlimited>]
[-MaxLocalHopCount **MaxHops**]
[-MaxLogonFailures **MaxLogonFailures**]
[-MaxMessageSize **MaxMessageSize**]
[-MaxProtocolErrors < **MaxErrors** | Unlimited>]
[-MaxRecipientsPerMessage **MaxRecipients**]
[-MessageRateLimit < **RateLimit** | Unlimited>]
[-MessageRateSource <User | IPAddress | Both>]
[-OrarEnabled <\$true | \$false>]

[-PermissionGroups <None | AnonymousUsers | ExchangeUsers | ExchangeServers | ExchangeLegacyServers | Partners | Custom >]
 [-PipeliningEnabled < \$true | \$false>]
 [-ProtocolLoggingLevel <None | Verbose>]
 [-RemoteIPRanges **IPRange1** , **IPRange2** , . . .]
 [-RequireEHLODomain <\$true | \$false>]
 [-RequireTLS < \$true | \$false>]
 [-Server **Server**]
 [-ServiceDiscoveryFqdn **ServiceFqdn**]
 [-SizeEnabled <Disabled | Enabled | EnabledWithoutValue>]
 [-SuppressXAnonymousTls < \$true | \$false>]
 [-TarpitInterval **TimeSpan**]
 [-TlsCertificateName "X509:<I> **Issuer** <S> **CommonName** "
 [-TlsDomainCapabilities **DomainName:Capability**]

Usage

New-ReceiveConnector -Name "Custom Receive Connector"
 -Usage "Custom" -Bindings "0.0.0.0:425"
 -Fqdn "mailserver85.tvpress.com"
 -RemoteIPRanges "0.0.0.0-255.255.255.255"
 -Server "CORPSVR127"
 -TransportRole HubTransport

Configuring Receive Connectors

To view all available Receive connectors for a server, select Mail Flow in the Navigation menu, and then select Receive Connectors. Next, on the Receive Connectors page, use the Select Server list to choose the server you want to work with. As shown in Figure 20-7, Receive connectors for the selected server are then listed by name, status, and role.

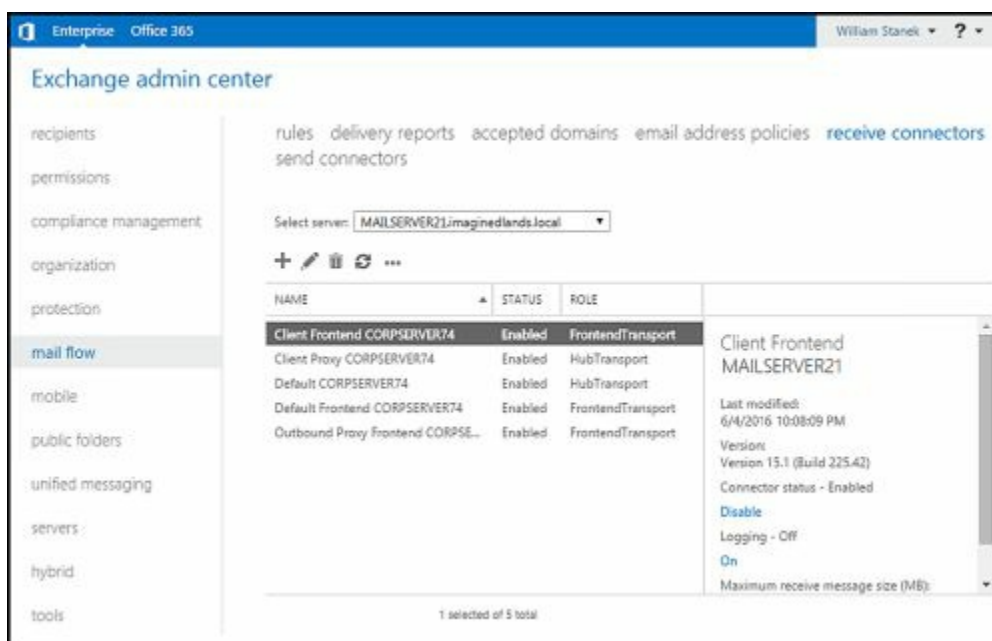



FIGURE 20-7 Working with receive connectors in Exchange Admin Center.

You can now:

- **Enable a connector** To enable a connector, select it, and then select Enable in the details pane.
- **Disable a connector** To disable a connector, select it, and then select Disable in the details pane.
- **Remove a connector** To remove a connector, select it in Exchange Admin Center, and then select Delete ().

To change a connector's properties, double-click the connector entry, and then use the Properties dialog box to manage the connector's properties, including protocol logging level, maximum receive size and maximum hop counts.



Client Frontend CORPSERVER74

▶ general
security
scoping

Protocol logging level:
 None
 Verbose

Maximum receive message size (MB):
35

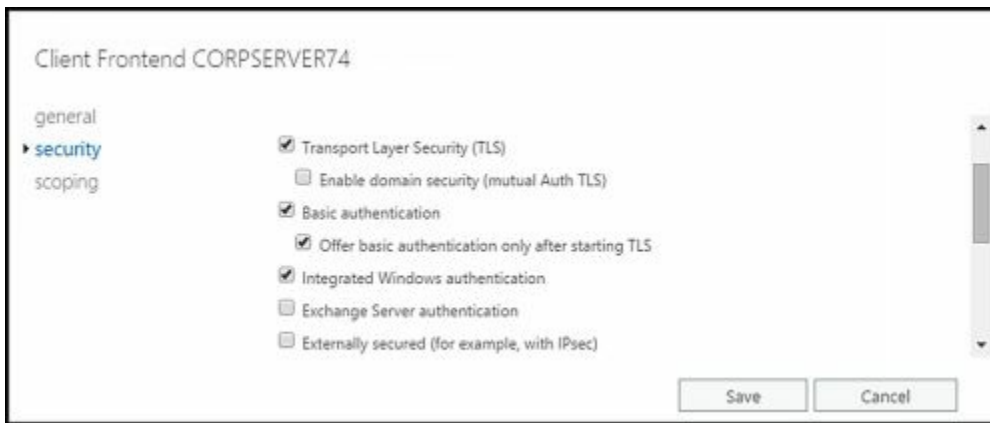
*Maximum local hop count:
12

*Maximum hop count:
60

Save Cancel

When configuring Receive connector properties, you can specify the security mechanisms that can be used for incoming connections on the Security page. Use any combination of the following:

- **Transport Layer Security** Allows encrypted authentications with TLS for servers with smart cards or X.509 certificates.
- **Enable Domain Security (Mutual Auth TLS)** When TLS is enabled, you can also enable domain security to require mutual authentication.
- **Basic Authentication** Allows basic authentication. With basic authentication, the user name and password specified are passed as base64-encoded text to the remote domain. Base64-encoding is cleartext and should not be confused with encryption.
- **Offer Basic Authentication Only After Starting TLS** Allows basic authentication only within an encrypted TLS session.
- **Integrated Windows Authentication** Allows secure authentication by using NT LAN Manager (NTLM) or Kerberos.
- **Exchange Server Authentication** Allows secure authentication for Exchange servers. With Exchange Server authentication, credentials are passed securely.
- **Externally Secured** Allows secure external authentication. With externally secured authentication, credentials are passed securely by using an external security protocol for which the server has been separately configured, such as IPsec.



Also when configuring Receive connector properties, you can specify the security group that is allowed to connect on the Permission Groups panel of the Security page. Use any combination of the following:

- **Anonymous Users** Allows unauthenticated, anonymous users to connect to the Receive connector.
- **Exchange Users** Allows connections by authenticated users who are valid recipients in the organization (Mailbox servers only).
- **Exchange Servers** Allows connections by authenticated servers that are members of the Exchange Server Administrator group.
- **Legacy Exchange Servers** Allows connections by authenticated servers that are members of the ExchangeLegacyInterop group.
- **Partners** Allows connections by authenticated servers that are members of partner domains, as listed on the TLS Receive Domain Secure list.



In Exchange Management Shell, you can view, update, or remove Receive connectors by using the `Get-ReceiveConnector`, `Set-ReceiveConnector`, or `Remove-ReceiveConnector` cmdlets, respectively. Listings 20-10 through 20-12 provide the syntax and usage. With `Get-ReceiveConnector`, you can return a list of all available Receive connectors if you don't specify an identity or server. If you want to see only the Receive connectors configured on a particular server, use the `-Server` parameter.

LISTING 20-10 `Get-ReceiveConnector` cmdlet syntax and usage

Syntax

```
Get-ReceiveConnector [-Identity Server\ConnectorIdentity ]
```

[-Server **Server**] [-DomainController **DCName**]

Usage

Get-ReceiveConnector

Get-ReceiveConnector -Identity "Corpsvr127\Imaginedlands.com Receive Connector"

Get-ReceiveConnector -Server "Corpsvr127"

LISTING 20-11 Set-ReceiveConnector cmdlet syntax and usage

Syntax

Set-ReceiveConnector -Identity **Identity**
[-AdvertiseClientSettings <\$true | \$false>]
[-AuthMechanism <None | Tls | Integrated | BasicAuth |
BasicAuthRequireTLS | ExchangeServer | ExternalAuthoritative>]
[-Banner **Banner**]
[-BareLineFeedRejectionEnabled <\$true | \$false>]
[-BinaryMimeEnabled <\$true | \$false>]
[-Bindings **Bindings**]
[-ChunkingEnabled <\$true | \$false >]
[-Comment **Comment**]
[-ConnectionInactivityTimeout **TimeSpan**]
[-ConnectionTimeout **TimeSpan**]
[-DefaultDomain **DefaultDomain**]
[-DeliveryStatusNotificationEnabled <\$true | \$false>]
[-DomainController **DCName**]
[-DomainSecureEnabled <\$true | \$false>]
[-EightBitMimeEnabled <\$true | \$false>]
[-EnableAuthGSSAPI <\$true | \$false>]
[-Enabled <\$true | \$false>]
[-EnhancedStatusCodesEnabled <\$true | \$false>]
[-ExtendedProtectionPolicy <none | allow | require>]
[-Fqdn **FQDN**]
[-LiveCredentialEnabled <\$true | \$false>]
[-LongAddressesEnabled <\$true | \$false>]
[-MaxAcknowledgementDelay **MaxDelay**]
[-MaxHeaderSize **MaxHeaderBytes**]
[-MaxHopCount **MaxHops**]
[-MaxInboundConnection < **MaxConn** | Unlimited>]
[-MaxInboundConnectionPercentagePerSource **MaxPercentage**]
[-MaxInboundConnectionPerSource < **MaxConnPerSource** | Unlimited>]
[-MaxLocalHopCount **MaxHops**]
[-MaxLogonFailures **MaxLogonFailures**]
[-MaxMessageSize **MaxMessageSize**]
[-MaxProtocolErrors < **MaxErrors** | Unlimited>]
[-MaxRecipientsPerMessage **MaxRecipients**]
[-MessageRateLimit < **RateLimit** | Unlimited>]

[-MessageRateSource <None | User | IPAddress | All>]
 [-Name **Name**]
 [-OrarEnabled <\$true | \$false>]
 [-PermissionGroups <None | AnonymousUsers | ExchangeUsers | ExchangeServers | ExchangeLegacyServers | Partners | Custom>]
 [-PipeliningEnabled < \$true | \$false>]
 [-ProtocolLoggingLevel <None | Verbose>]
 [-RemoteIPRanges **IPRange1** , **IPRange2** , . . .]
 [-RequireEHLODomain <\$true | \$false>]
 [-RequireTLS < \$true | \$false>]
 [-ServiceDiscoveryFqdn **ServiceFqdn**]
 [-SizeEnabled <Disabled | Enabled | EnabledWithoutValue>]
 [-SuppressXAnonymousTls < \$true | \$false>]
 [-TarpitInterval **TimeSpan**]
 [-TlsCertificateName "X509:<I> **Issuer** <S> **CommonName** "
 [-TlsDomainCapabilities **DomainName:Capability**]
 [-TransportRole <None | Cafe | Mailbox | ClientAccess | UnifiedMessaging | HubTransport | Edge | All | Monitoring | CentralAdmin | CentralAdminDatabase | DomainController | WindowsDeploymentServer | ProvisionedServer | LanguagePacks | FrontendTransport | CafeArray | FfoWebService | OSP | ARR | ManagementFrontEnd | ManagementBackEnd | SCOM>]

Usage

```
Set-ReceiveConnector -Identity "Corpsvr127\Custom Receive Connector"
-Bindings "0.0.0.0:425"
-Fqdn "mailserver85.tvpress.com"
-RemoteIPRanges "0.0.0.0-255.255.255.255"
```

Listing 20-12 Remove-ReceiveConnector cmdlet syntax and usage

Syntax

```
Remove-ReceiveConnector -Identity ConnectorIdentity
[-Confirm <$true | $false >]
[-DomainController DCName ]
```

Usage

```
Remove-ReceiveConnector -Identity "CorpSvr127\Imaginedlands.com Receive Connector"
```

Creating Connectors with Exchange Online

Exchange Online uses Inbound and Outbound connectors, rather than Receive and Send connectors. Inbound connectors control mail flowing from the Internet, a partner, or a specific server. Outbound connectors control the flow of mail sent by recipients in the organization. When mailbox users in the online organization are sending mail, you can use Outbound connectors to direct messages to a server that applies additional processing before delivering the mail to its destination.

When you run the Hybrid Configuration Wizard to create a hybrid organization that

combines an on-premises Exchange organization with an online Exchange organization, a Send connector is created automatically in the on-premises Exchange organization to route mail to Exchange Online and a Receive connector is created automatically to receive mail from Exchange Online. Similarly, an Outbound connector is created automatically in the online Exchange organization to route mail to on-premises Exchange, and an Inbound connector is created automatically to receive mail from on-premises Exchange.

The automatically created Inbound and Outbound connectors have the connector type set as On-Premises. To view and manage inbound or outbound connectors, access Exchange Admin Center for Exchange Online. Next, select Mail Flow in the Navigation menu, and then select Connectors, as shown in Figure 20-8.

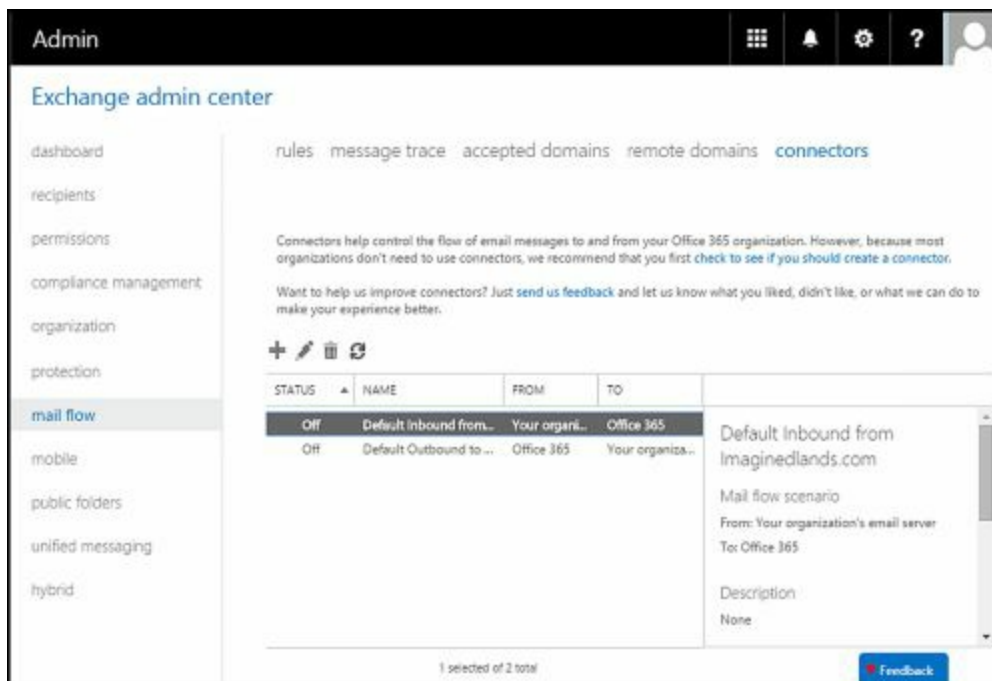
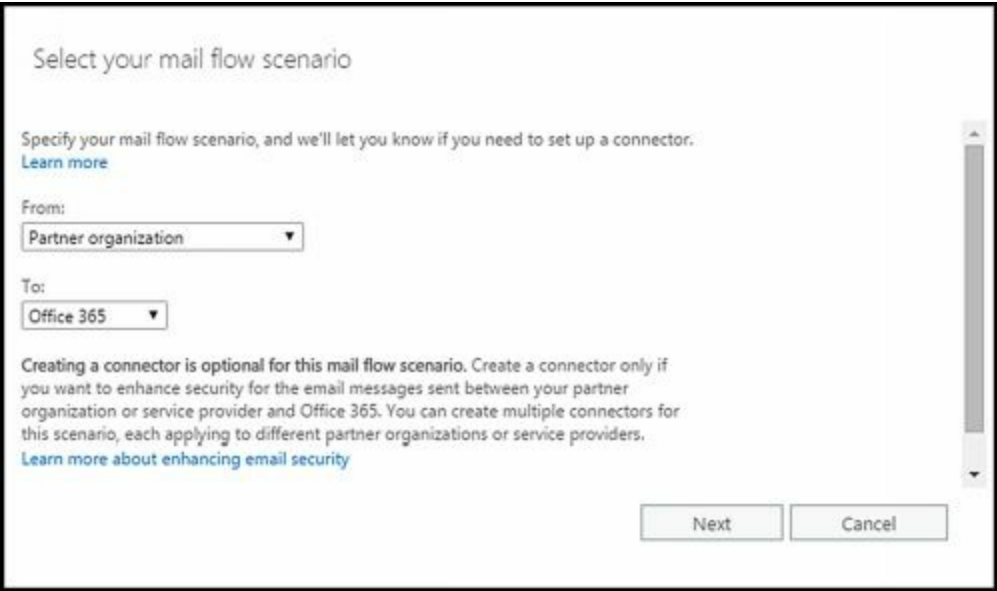


FIGURE 20-8 Viewing connectors in Exchange Online.

You can create additional Inbound and Outbound connectors to control mail flow from and to trusted partners. These additional connectors have the connector type set as Partner, rather than On-Premises. By default, connectors use opportunistic TLS for connection security. This means connectors try to use TLS security for connections but if TLS cannot be used, they establish a standard SMTP connection instead.

To create Inbound or Outbound connectors, click Add (**+**). Use the selection lists provided to specify where messages are being routed from and to, such as from Internet to Office 365 or from a partner organization to Office 365. Next, use the options provided to configure the connectors much as you would configure Send and Receive connectors.



Select your mail flow scenario

Specify your mail flow scenario, and we'll let you know if you need to set up a connector.
[Learn more](#)

From:
Partner organization ▼

To:
Office 365 ▼

Creating a connector is optional for this mail flow scenario. Create a connector only if you want to enhance security for the email messages sent between your partner organization or service provider and Office 365. You can create multiple connectors for this scenario, each applying to different partner organizations or service providers.
[Learn more about enhancing email security](#)

Next Cancel

You also can connect to Exchange Online in Windows PowerShell and then use the `New-InboundConnector` or `New-OutboundConnector` cmdlet to create a connector. Each connector type has corresponding `Set`, `Get`, and `Remove` cmdlets as well. These are `Set-InboundConnector`, `Get-InboundConnector`, and `Remove-InboundConnector` as well as `Set-OutboundConnector`, `Get-OutboundConnector`, and `Remove-OutboundConnector`.

Chapter 21. Configuring Transport Services

You can configure your Microsoft Exchange Server 2016 organization with only Mailbox servers for message routing and delivery, or you can configure it with Mailbox servers and Edge Transport servers. When you use only Mailbox servers, these servers are responsible for:

- Messaging routing and delivery within the organization.
- Receiving messages from outside the organization and delivering them to Mailbox servers within the organization.
- Receiving messages from Mailbox servers within the organization and routing them to destinations outside the organization.

When you use both Mailbox servers and Edge Transport servers, message routing and delivery works like this:

- Mailbox servers handle message routing and delivery within the organization.
- Edge Transport servers receive messages from outside the organization and route them to Mailbox servers within the organization which, in turn, deliver them to other Mailbox servers if necessary.
- Mailbox servers receive messages from Mailbox servers within the organization and route them to Edge Transport servers, which, in turn, route them to destinations outside the organization.

NOTE In a mixed environment where Exchange 2010 Hub Transports are deployed, the Hub Transport servers work with Mailbox servers and Edge Transport servers to route and deliver messages.

When you use Edge Transport servers in a hybrid deployment, your Edge Transport servers can be configured to handle communications between on-premises Exchange and Exchange Online. Here, the Edge Transport servers act as relays between your internal Exchange servers and Exchange Online, as long as the Edge Transport servers are externally accessible from the Internet on port 25. Additionally, at this time, only Edge Transport servers running Exchange 2010 Service Pack 2 or later support hybrid deployments.

The primary mail protocol used by Exchange Server 2016 is Simple Mail Transfer Protocol (SMTP). This chapter discusses how transport servers use SMTP for routing and delivery, as well as how you can view and manage transport server configurations.

REAL WORLD Microsoft recommends that you install the Edge Transport server role on a computer that is not part of the internal Active Directory domain. The server can, however, be part of an external Active Directory domain, which isolates the computer and is the most secure implementation. Although you can install the Edge Transport server on a domain-joined computer, the Edge Transport server role will always use Active Directory Lightweight Directory Services (AD

LDS) to store recipient and configuration information for the Edge stack, and the underlying Windows stack will use Active Directory Domain Services (AD DS). To send and receive messages from your organization to the Internet, Edge Transport servers use Send connectors and Receive connectors.

Prior to installing the Edge Transport role, you need to set the Domain Name System (DNS) suffix for the server and install the AD LDS role. Generally, you'll want to use a DNS suffix for your organization's primary domain. To install the AD LDS role, use the Add Roles Wizard in the Server Manager. Accept the default settings during installation with one exception: you do not need to create an application partition because AD LDS will be configured for the Edge Transport server role when you install the role, and the required application partition will also be created at that time.

Optimizing Transport Limits

Exchange Server 2016 automatically places receive size, send size, and other limits on messages being routed through an Exchange organization. The limits you can control include:

- Message header limits control the total size of all message header fields in a message. Header limits primarily apply to Receive connectors, although they also apply to messages in the pickup directory used by the Transport service. Header fields are plain text, and so the size of the header is determined by the total number of header fields and characters in each header field. Each character of text is 1 byte.
- Message receive size limits control the total size of messages that can be received, which includes the message header, message body, and any attachments. Exchange uses a custom message header (X-MS-Exchange-Organization-OriginalSize:) to record the original size of a message when it enters the Exchange organization. Although content conversion, encoding, and agent processing can change the size of the message, Exchange uses the lower value of the current or original message size to determine whether the limit applies.
- Message send size limits control the total size of messages that can be sent, which includes the message header, message body, and any attachments.
- Attachment size limits control the maximum size of each individual attachment within a message.
- Recipient limits control the total number of message recipients, with an unexpanded distribution group counted as a single recipient. When a message is composed, recipients are listed in the To:, Cc:, and Bcc: header fields. When a message is submitted for delivery, these recipients are converted into Rcpt To: entries in the message.

NOTE Unlike other limits, exceeding a recipient limit doesn't automatically mean a message will be rejected. The message may be accepted for the first N recipients and then resent by the SMTP server in groups of N recipients until the message is delivered to all recipients.

A message that exceeds any applicable limit is rejected and a non-delivery report is issued to the sender with an error code, status, and description. Transport limits are configured for the organization as a whole, for individual send and receive connectors, for specific servers, for specific users, and for specific Active Directory site links.

As part of your planning for message size limits, you need to consider that base64 encoding will be applied to attachments and any binary data in messages. Base64 encoding increases the size of the attachments and the binary data by approximately 33 percent and in this way increases the total size of a message. Thus, attachments with a total original size of 27 MB could cause a message to exceed a send or receive limit of 35 MB.

Setting Organizational Transport Limits

Organizational transport limits apply to all transport servers in the organization, which includes Exchange 2016 Mailbox servers, Exchange 2010 Hub Transport servers and Exchange 2007 Hub Transport servers. By default, the maximum message size that can be received or sent by recipients in the organization is 10,240 KB and messages can have no more than 500 recipients.

You can view or change the default limits for the Exchange organization by completing the following steps:

1. In the Exchange Admin Center, select Mail Flow in the Navigation menu, and then select either Receive Connectors or Send Connectors.
2. In the main pane, select the More button (**⋮**), and then click Organization Transport Settings. In the Organization Transport Settings dialog box, the Limits page is selected by default, as shown in Figure 21-1.



FIGURE 21-1 Set transport limits for the Exchange organization.

3. To set a maximum number of recipients limit, type the desired limit in the Maximum Number Of Recipients combo box. The valid input range is 0 to 2,147,483,647. If you use a value of 0, no limit is imposed on the number of recipients in a message. Note that Exchange handles an unexpanded distribution group as one recipient.
4. To set a maximum receive size limit, type the desired receive limit in the related combobox. The valid input range is 0 to 2047.999 MB. If you use a value of 0 or select Unlimited in the dropdown list, no limit is imposed on the message size that can be received by recipients in the organization.
5. To set a maximum send size limit, type the desired send limit in the related combobox. The valid input range is 0 to 2047.999 MB. If you use a value of 0 or select Unlimited in the dropdown list, no limit is imposed on the message size that can be sent by senders in the organization.
6. Click Save to apply your settings.

In Exchange Management Shell, you assign the desired transport limits by using the Set-

TransportConfig cmdlet, as shown in Listing 21-1. The `-MaxReceiveSize` and `-MaxSendSize` parameters set the maximum receive size and maximum send size, respectively. The `-MaxRecipientEnvelopeLimit` parameter sets the maximum number of recipients in a message. When you use the `-MaxReceiveSize` and `-MaxSendSize` parameters, you must specify the units for values by using KB for kilobytes, MB for megabytes, or GB for gigabytes. Your changes are made at the organization level and apply to the entire Exchange Server 2016 organization.

LISTING 21-1 Setting transport limits

Syntax

```
Set-TransportConfig [-Identity OrgId ] [-DomainController DCName ]  
[-MaxReceiveSize <' MaxSize ' | 'Unlimited'>]  
[-MaxSendSize <' MaxSize ' | 'Unlimited'>]  
[-MaxRecipientEnvelopeLimit <' MaxRecipients ' | 'Unlimited'>]
```

Usage

```
Set-TransportConfig -MaxReceiveSize '15MB' -MaxSendSize '15MB'  
-MaxRecipientEnvelopeLimit '1000'
```

You can control the maximum message size and maximum attachment size for all Mailbox servers in the organization by using transport rules. To do this, use the `-MessageSizeOver` and `-AttachmentSizeOver` parameters of `New-TransportRule` or `Set-TransportRule`.

Setting Connector Transport Limits

The transport limits of a connector apply to any message that uses a specified connector for message delivery. Exchange 2016 automatically sets transport limits on Send and Receive connectors. Most connectors have a maximum message size limit of 35 MB by default. The exceptions are the Default Frontend and Outbound Proxy Frontend Receive connectors, which have a 36 MB limit by default.

You can view the current maximum message size limits for all send connectors by entering the following command in Exchange Management Shell:

```
get-sendconnector | fl name, maxmessagesize
```

To view the current maximum size of all receive connectors, enter:

```
get-receiveconnector | fl name, maxmessagesize
```

You can modify the default maximum message size limit by using the `-MaxMessageSize` parameter of the `New-ReceiveConnector`, `Set-ReceiveConnector`, `New-SendConnector`, and `Set-SendConnector` cmdlets.

Receive connectors also have default limits on the maximum number of recipients and the maximum header size. Most of the default Receive connectors have a limit of 200

recipients by default. The exception is the Default Receive connector which has a limit of 5,000 recipients by default.

The default Receive connectors and any other Receive connectors you create automatically have a 128 KB maximum header size limit. Although Exchange adds headers to messages during content conversion, encoding, and agent processing, the number of recipients in a message is the most common reason a message exceeds the maximum header size limit. Each character in a recipient's name and email address counts against the limit. If a message is rejected because it exceeds the maximum header size limit, the sender should receive a non-delivery report. This non-delivery report may contain an error status code of 4.4.7, which can help you identify the problem as relating to the maximum header size limit.

In the shell, you can view the current recipient and header size limits for all receive connectors by entering:

```
get-receiveconnector |fl name, maxheadersize, maxrecipientspermessage
```

You can modify the recipient and header limits by using the `-MaxRecipientsPerMessage` and `-MaxHeaderSize` parameters of the `New-ReceiveConnector` and `Set-ReceiveConnector` cmdlets.

Setting Server Transport Limits

The transport limits of a server apply to any message processed by the server. If a user's mailbox is on a particular Mailbox server, the maximum header size and maximum number of recipient limits for the pickup directory apply. You can configure these limits on a per-server basis as discussed in Chapter 22, "Maintaining Mail Flow" in the "Configuring Messaging Limits for the Pickup Directory" section.

Per-server transport limits also apply to the front-end transport services. The maximum message size for Outlook Web App is 33 MB. Exchange ActiveSync and Exchange Web Services have maximum message size limits of 10 MB and 64 MB respectively. To change these values, you must edit the appropriate `web.config` configuration file on each Mailbox server. The configuration files are formatted with XML and can be edited in any standard text editor, including Notepad.exe.

IMPORTANT Before you make any changes, you might want to create a copy of each of the original configuration files. In Notepad, you can use the Find feature on the Edit menu to search. As the default search starts at the current position, make sure you start your searches at the top of the document. One way to ensure you are at the top of the document is to press `Ctrl+Home` while working in Notepad.

Setting Exchange Activesync Limits

The `%ExchangeInstallpath%` variable is an environment variable set when you installed Exchange server. You'll find the `web.config` file for Exchange ActiveSync in the `%ExchangeInstallpath%\ClientAccess\Sync` folder. In this `web.config` file, the

MaxDocumentDataSize key sets the maximum size of data that can be received by the ActiveSync protocol, and the MaxRequestLength value sets the maximum size of data that can be received from an ActiveSync client.

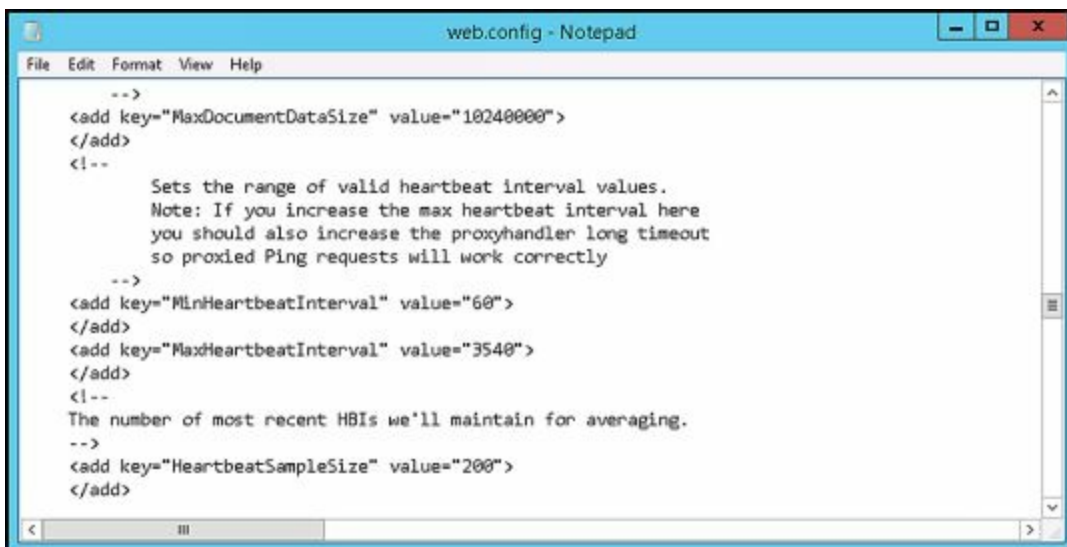
You can open the configuration file for editing in Notepad by entering the following command at a command prompt:

```
Notepad %ExchangeInstallpath%\ClientAccess\Sync\web.config
```

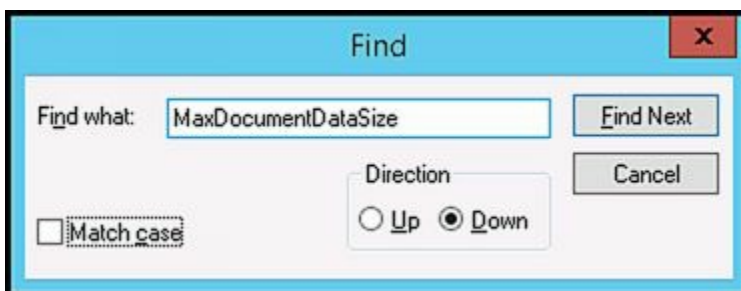
Or entering the following at the PowerShell prompt:

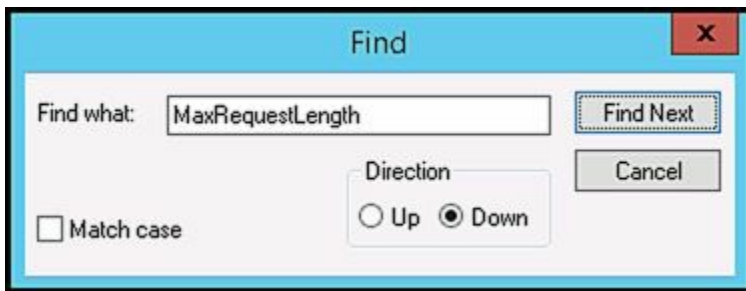
```
Notepad $env:ExchangeInstallpath\ClientAccess\Sync\web.config
```

REAL WORLD If you're using Exchange Management Shell rather than a standard PowerShell prompt, keep in mind Exchange Management Shell does not run in elevated, administrator mode by default because your login credentials are used to create an implicit remoting session. Although you can run Exchange Management Shell in administrator mode, a new session for remoting won't be implicitly established until you run the first Exchange command.



After you open the web.config file, search for MaxDocumentDataSize, and then set the related value to the desired maximum size in kilobytes (KB). Next, search for MaxRequestLength to set the related value to the desired maximum size in bytes.





The related entries are:

```
<add key="MaxDocumentDataSize" value="10240000">
... </add>
<httpRuntime maxRequestLength="10240" />
```

When you are finished making changes, save and close the configuration file. Keep in mind that when you save the changes to the configuration file, the related web application is restarted automatically.

Confirm that Exchange ActiveSync is working as expected by entering **Test-ActiveSyncConnectivity** at the shell prompt. If there's a problem with Exchange ActiveSync, check your edits or restore the original configuration file.

Setting Exchange Web Services Limits

You'll find the web.config file for Exchange Web Services in the %ExchangeInstallpath%\ClientAccess\exchweb\ews folder. In this web.config file, the MaxAllowedContentLength value sets the maximum size of HTTP content requests and the MaxReceivedMessageSize value sets the maximum size of messages that can be accepted by Exchange Web Services.

You can open the configuration file for editing in Notepad by entering the following command at a command prompt:

```
Notepad %ExchangeInstallpath%\ClientAccess\exchweb\ews\web.config
```

Or entering the following at the shell prompt:

```
Notepad $env:ExchangeInstallpath\ClientAccess\exchweb\ews\web.config
```




```
web.config - Notepad
File Edit Format View Help
<customBinding>
  <binding name="EWSAnonymousHttpsBinding">
    <EWSMessageEncoderSoap11Element />
    <!-- Since we are using a transfer mode of "streamed" maxBufferSize can be used to limit the size of
         the SOAP header section since that is the only thing that is buffered when receiving a stream
         See http://kenryw.com/indigo/78 -->
    <httpsTransport authenticationScheme="Anonymous" maxReceivedMessageSize="67108864" maxBufferSize="163840" />
  </binding>
  <binding name="EWSAnonymousHttpBinding">
    <EWSMessageEncoderSoap11Element />
    <!-- Since we are using a transfer mode of "streamed" maxBufferSize can be used to limit the size of
         the SOAP header section since that is the only thing that is buffered when receiving a stream
         See http://kenryw.com/indigo/78 -->
    <httpTransport authenticationScheme="Anonymous" maxReceivedMessageSize="67108864" maxBufferSize="163840" />
  </binding>
  <binding name="EWSBasicHttpsBinding">
    <EWSMessageEncoderSoap11Element />
    <!-- Since we are using a transfer mode of "streamed" maxBufferSize can be used to limit the size of
         the SOAP header section since that is the only thing that is buffered when receiving a stream
         See http://kenryw.com/indigo/78 -->
    <httpsTransport authenticationScheme="Basic" maxReceivedMessageSize="67108864" maxBufferSize="163840" />
  </binding>
  <binding name="EWSBasicHttpBinding">
    <EWSMessageEncoderSoap11Element />
  </binding>
</customBinding>
```

After you open the web.config file, search for each occurrence of MaxReceivedMessageSize to set the related value to the desired maximum size in bytes. You must set a MaxReceivedMessageSize value for each HTTP and HTTPS binding and authentication combination.



IMPORTANT Although there are 16 entries for MaxReceivedMessageSize that you want to edit in total, you don't want to modify the two entries for UM bindings.

Next, search for MaxAllowedContentLength and then set the related value to the desired maximum size in bytes. When you're finished making changes, save and close the configuration file. Keep in mind that when you save the changes to the configuration file, the related web application is restarted automatically.



Confirm that Exchange Web Services are working as expected by entering **Test-WebServicesConnectivity** at the shell prompt. If a problem occurs with Exchange Web Services, check your edits or restore the original configuration file.

Setting Outlook Web App Limits

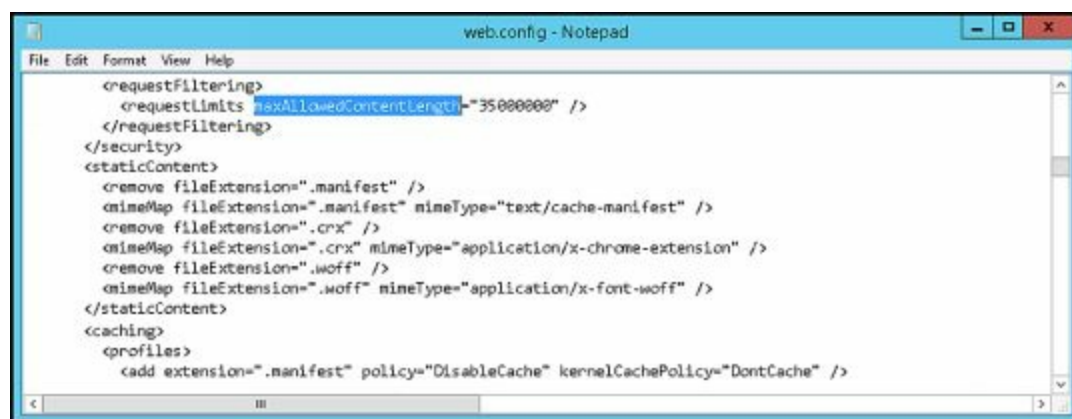
You'll find the web.config file for Outlook Web App in the %ExchangeInstallpath%\ClientAccess\Owa folder. In this web.config file, the MaxAllowedContentLength value key sets the maximum size of HTTP content requests, MaxReceivedMessageSize value sets the maximum size of messages that can be accepted by Outlook Web App and MaxRequestLength value sets the maximum size of data that can be received from an Outlook Web App client.

You can open the configuration file for editing in Notepad by entering the following command at a command prompt:

```
Notepad %ExchangeInstallpath%\ClientAccess\Owa\web.config
```

Or entering the following at the shell prompt:

```
Notepad $env:ExchangeInstallpath\ClientAccess\Owa\web.config
```



After you open the web.config file, search for MaxAllowedContentLength, and then set the related value to the desired maximum size in bytes. Next, search for MaxReceivedMessageSize, and then set the related value to the desired maximum size in bytes. There are two entries for MaxReceivedMessageSize: one for HTTP and one for HTTPS. Finally, search for MaxRequestLength to set the related value to the desired maximum size in kilobytes. The related entries are:

```
<requestLimits maxAllowedContentLength="35000000" />
...
<binding name="httpsBinding" maxReceivedMessageSize="35000000">
...
<binding name="httpBinding" maxReceivedMessageSize="35000000">
...
<httpRuntime maxUrlLength="500" maxRequestLength="35000"
requestValidationMode="2.0" enableVersionHeader="false" />
```

When you are finished making changes, save and close the configuration file. Keep in mind that when you save the changes to the configuration file, the related web application is restarted automatically. Confirm that Outlook Web App is working as expected by entering **Test-OwaConnectivity** at the shell prompt. If a problem with Outlook Web App occurs, check your edits or restore the original configuration file.

REAL WORLD By default, IIS uses overlapping recycling of worker processes when restarting applications and application pools. With overlapping recycling,

new worker processes are started to accept new requests from HTTP.sys while current worker processes are marked for recycling but continue handling existing requests. When all existing requests are handled, the original worker processes shut down.

Managing Message Transport

After you install Mailbox servers running Exchange Server 2016, you need to finalize the configuration of Transport services by creating and configuring a postmaster mailbox and performing any other necessary tasks. For Exchange organizations with only Mailbox servers, you should optimize anti-spam features. For Exchange organizations with Edge Transport servers, you need to subscribe the Edge Transport servers to your Exchange organization.

Configuring the Postmaster Address and Mailbox

Every organization that sends and receives mail should have a postmaster address. This is the email address listed on nondelivery reports and other delivery status notification reports created by Exchange Server. The postmaster address is not set by default; therefore, you must manually set it.

To view your Exchange organization's postmaster address, enter the following command at Exchange Management Shell prompt:

```
Get-TransportConfig | Format-List Name,ExternalPostMasterAddress
```

This command lists the postmaster address for the organization, as shown in this sample output:


```
Name:                Transport Settings
ExternalPostmasterAddress : postmaster@tvpress.com
```

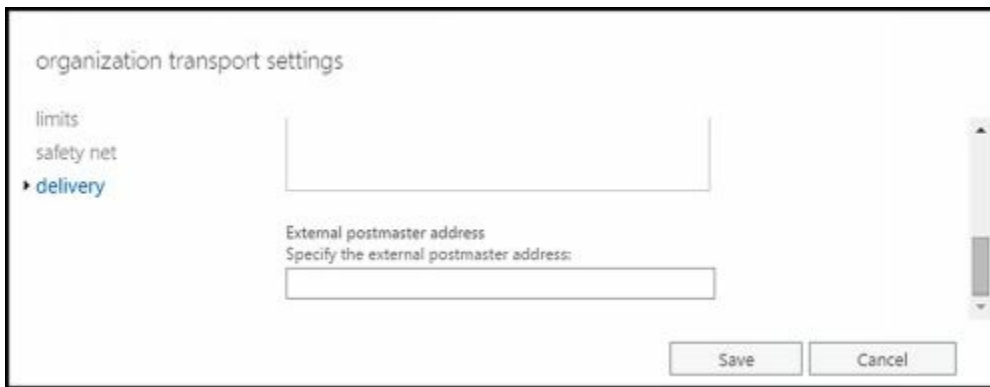
If you don't set the postmaster address, the address typically is set to \$null, except when you have an Edge Transport server that hasn't been through the Edge Sync process. To change the postmaster address, you can use the `-ExternalPostMasterAddress` parameter of the `Set-TransportServer` cmdlet, as shown in this example:

```
Set-TransportConfig -ExternalPostMasterAddress "nondelivery@tvpress.com"
```

If you want the postmaster address to be able to receive mail, you must either create a mailbox and associate it with the postmaster address or assign the postmaster address as a secondary email address for an existing mailbox.

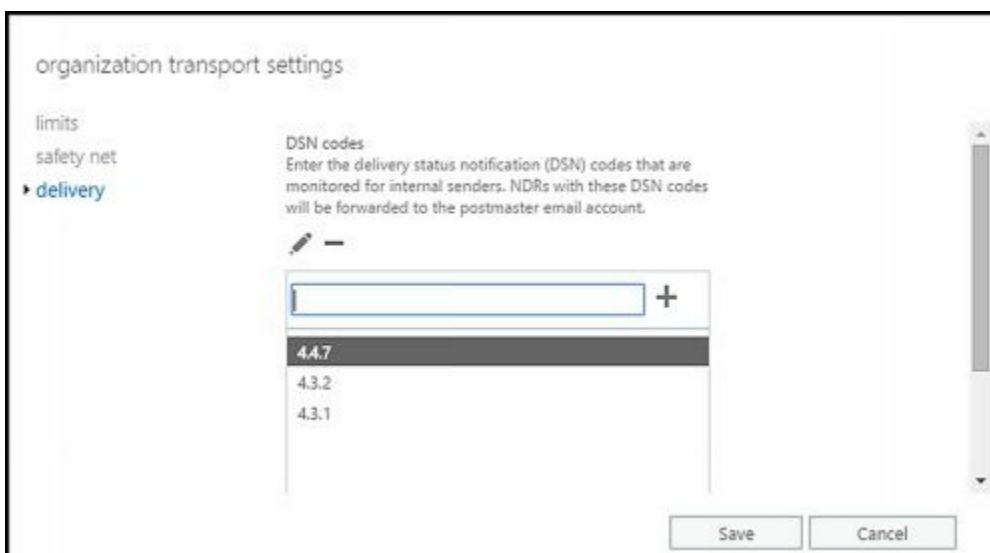
You also can view or change the organization's postmaster address by completing the following steps:

1. In the Exchange Admin Center, select Mail Flow in the Navigation menu, and then select either Receive Connectors or Send Connectors.
2. In the main pane, select the More button (), and then click Organization Transport Settings to display the Organization Transport Settings dialog box.
3. On the Delivery page, the current postmaster email address is listed (if any). If you want to change the postmaster address, enter the address you want to use, and then click Save.



On the Delivery page, you also can specify the delivery status notification (DSN) codes that should be monitored. The postmaster receives a copy of any non-delivery reports delivered to internal senders with these codes. Codes you may want to have monitored include:

- **4.3.1** Issued when there are insufficient resources on the Mailbox server, usually as a result of a resource problem. Note that the report may state an out-of-memory error when the actual error that occurred was caused by a full disk.
- **4.3.2** Issued when the system is not accepting network messages often caused by a frozen queue. To resolve the problem, unfreeze the queue.
- **4.4.7** Issued when a message expires before it can be relayed or delivered, typically occurring as a result of a time-out during communication with a remote server. It also can indicate a message header limit has been reached, so the sender may need to reduce the number of recipients in the message.
- **5.3.5** Issued when a server is improperly configured, specifically when the server is configured to loop mail back to itself. To resolve the problem, check the server's connectors for loops.
- **5.4.6** Issued when a routing loop is detected, specifically when the delivery of a message generates another message and that message then generates another message, and so forth. If the message generating loop continues more than 20 times, this error is issued. To resolve the error, check the mailbox rules associated with the recipients and senders to determine how automatic message forwarding is configured.



Configuring Shadow Redundancy

Shadow redundancy ensures that messages are protected from loss the entire time they are in transit by creating a copy of a message and retaining this copy while a message is in transit. If any transport server along the route fails to report a successful delivery, Exchange resubmits the message for delivery to ensure that the message continues through to its destination.

By default, shadow redundancy is enabled in the Transport service on all Mailbox servers. Exchange 2016 makes a redundant copy of any message it receives before acknowledging receipt. This approach ensures the message will be delivered even if the receiving server were to shut down immediately after acknowledging receipt of a message. Prior to this approach, a message could possibly be lost if the receiving server were to shut down after acknowledging receipt of a message but before creating a copy of the message.

Thanks to shadow redundancy, as long as you have multiple transport servers (and multiple Edge Transports if you've deployed Edge Transport servers), you can remove any transport server that fails and not have to worry about emptying its queues or losing messages. You also can upgrade or replace a Mailbox or Edge Transport server at any time without the risk of losing messages. If you have a single Mailbox server, you should drain all SMTP queues on the server before performing maintenance. The same is true if you have a single Edge Transport. This ensures that there is no risk of message loss, even without shadow redundancy. Keep in mind that if you have a single transport server, and it fails and must be replaced, you've likely lost data if you can't restore the mail.que file.

IMPORTANT Shadow redundancy requires multiple servers. Your Mailbox servers can be standalone servers or they can be part of a database availability group. However, with standalone servers, each Active Directory site with Mailbox servers must have two or more standalone servers. Although there must be multiple members of a database availability group for shadow redundancy to work, the members of that group can be in different Active Directory sites.

When you work with shadow redundancy, a key concept to understand is that the primary transport server has ownership of the messages in its shadow queue. The first primary owner is always the server on which the message originates. As the message travels through the transport pipeline, different transport servers may become the primary owner of a message. In addition, if a primary owner fails, another server can take over as the primary.

Shadow redundancy is implemented according to high availability transport (HAT) boundaries in the organization. Each Active Directory site with Mailbox servers in the organization is a HAT boundary, as is each Database Availability group in the organization. Within a HAT boundary, two copies of a message are always in transit: the original and the redundant copy.

It's important to point out that the original copy and the redundant copy exist on different servers. When a Mailbox server receives a message, it makes a redundant copy of the message on another Mailbox server in the HAT boundary before acknowledging receipt of the message. With database availability groups, the Transport service prefers creating the redundant copy in a remote site to ensure site resilience.

The basic process works like this:

1. The primary server transmits a copy of the message to the Transport service on another Mailbox server, and the Transport service on the other Mailbox server acknowledges that the copy of the message was created successfully. The copy of the message is the shadow message, and the Mailbox server that holds it is the shadow server for the primary server. The message exists in a shadow queue on the shadow server.
2. After the primary server receives acknowledgement from the shadow server, the primary server acknowledges the receipt of the primary message to the original SMTP server in the original SMTP session, and the SMTP session is closed.
3. The primary server transmits the message. If the primary server transmits the message outside the HAT boundary and the receiving SMTP server acknowledges successful receipt of the message, the primary server moves the primary message into its Safety Net queue. Otherwise, if the ultimate destination for the message is within the HAT boundary, the primary message is moved into the Safety Net queue when the message is accepted by the Transport service on a Mailbox server that holds the ultimate destination for the message.
4. The shadow server moves the shadow message to its Safety Net queue.

This process is complex and can be difficult to understand, so let's take another look at the process. Step-by-step, the process works like this:

1. An SMTP server transmits a message to the Transport service on a Mailbox server in the Exchange organization. The receiving Mailbox server becomes the primary server for the message, and the original message is the primary message.
2. While the original SMTP session with the SMTP server is still active, the Transport service on the primary server opens a new, simultaneous SMTP session with the Transport service on another Mailbox server in the HAT boundary.
3. The primary server transmits a copy of the message to the Transport service on the other Mailbox server. The copy of the message is the shadow message, and the Mailbox server that holds it is the shadow server for the primary server. The message exists in a shadow queue on the shadow server.
4. After the primary server receives an acknowledgement from the shadow server that confirms the copy of the message was created, the primary server acknowledges the receipt of the primary message to the original SMTP server in the original SMTP session, and the SMTP session is closed.
5. The primary server transmits the message. The primary server and the shadow

server stay in contact with each other to track the progress of the message.

6. When the primary server successfully transmits the message and the receiving SMTP server acknowledges successful receipt of the message, the primary server updates the discard status of the message to show delivery is complete and relays this to the shadow server.
7. The shadow server moves the shadow message from the shadow queue to its Safety Net queue.

In Exchange Management Shell, you configure shadow redundancy for the on-premises Exchange organization by using the `Set-TransportConfig` cmdlet, as shown in Listing 21-2. The related parameters are used as follows:

- **MaxDumpsterTime** Only used by Hub Transport servers in a coexistence scenario. Specifies the maximum amount of time that a delivered message will remain in the transport dumpster for possible resubmission. The default is seven days.
- **MaxRetriesForLocalSiteShadow** When member servers in a database availability group span multiple Active Directory sites and `ShadowMessagePreferenceSetting` is configured to prefer remote sites, you can use this option to control how many times the primary server tries to create the shadow copy on a server in the local site after failing to create the copy in a remote site. By default, this option is set to 2. If the preference is for `LocalOnly`, this option controls the number of times the primary server tries to create the shadow copy on a server in the local site before failing and rejecting the message with a transient error.
- **MaxRetriesForRemoteSiteShadow** When member servers in a database availability group span multiple Active Directory sites and `ShadowMessagePreferenceSetting` is configured to prefer remote sites, you can use this option to control how many times the primary server tries to create the shadow copy on a server in a remote site before trying to create the shadow copy on a server in the local site. By default, this option is set to 4. If the preference is for `RemoteOnly`, this option controls the number of times the primary server tries to create the shadow copy on a server in a remote site before failing and rejecting the message with a transient error.
- **RejectMessageOnShadowFailure** Determines whether a primary message can be accepted or acknowledged without a shadow copy being created first. This option is disabled by default. If you enable this option and a shadow copy cannot be created, the primary message will be rejected with a transient error. Enable this option only when you must ensure a shadow copy of a message is always created and multiple Mailbox servers exist in each HAT boundary.
- **ShadowHeartbeatFrequency** Sets the amount of time a transport server waits before establishing a connection to the primary server to check the discard status of shadow messages. The default value is two minutes. Set this value according to the size of your Exchange implementation, the level of messaging traffic, and the relative latency on the network. For example, in a large global organization where transport servers handle an extremely high volume of messages, you might want to set a longer time interval, although the default may suffice for a smaller organization.

- **ShadowMessageAutoDiscardInterval** Sets the amount of time a server retains discard events for successfully delivered shadow messages. Primary servers queue discard events until they are checked by the shadow server or until the discard interval has elapsed, whichever comes first. The default value is two days. Set the value according to the size of your Exchange implementation, the level of messaging traffic, and the relative reliability of your network. For example, in a large global organization where transport servers handle an extremely high volume of messages on a highly reliable network, you might want to set a shorter discard interval, whereas the default may suffice for a smaller organization.
- **ShadowMessagePreferenceSetting** When member servers in a database availability group span multiple Active Directory sites, you can use this option to control remote site preferences. By default, this option is set to PreferRemote. Here, the primary server attempts to create a shadow copy on a server in a remote site. If this fails, the primary server attempts to create a shadow copy on a server in the local site. Alternatively, you can specify that the copy should only be made in the local site or only in a remote site. To do this, set the value to LocalOnly or RemoteOnly respectively.
- **ShadowRedundancyEnabled** Enables or disables shadow redundancy. If you don't use shadow redundancy, you can use this parameter to disable the feature. Ideally, you'd only disable the feature temporarily or in situations in which you have a single Exchange server implementation and are experiencing problems related to this feature. By default, shadow redundancy is enabled.
- **ShadowResubmitTimeSpan** Specifies how long a shadow server waits before deciding that the primary server has failed and assumes ownership of messages in the shadow queue for that server. The default value is three hours. Set this value according to the size of your Exchange implementation and the relative amount of latency on the network. For example, a large global organization might want to set a longer time span, whereas the default may suffice for a smaller organization.

LISTING 21-2 Setting shadow queue options

Syntax

```
Set-TransportConfig [-Identity OrgId ] [-DomainController DCName ]
[-MaxDumpsterTime < TimeSpan >]
[-MaxRetriesForLocalSiteShadow RetryCount ]
[-MaxRetriesForRemoteSiteShadow RetryCount ]
[-RejectMessageOnShadowFailure <$true | $false>]
[-SafetyNetHoldTime < TimeSpan >]
[-ShadowHeartbeatFrequency < TimeSpan >]
[-ShadowMessageAutoDiscardInterval < TimeSpan >]
[-ShadowMessagePreferenceSetting <PreferRemote | LocalOnly | RemoteOnly>]
[-ShadowRedundancyEnabled <$true | $false>]
[-ShadowResubmitTimeSpan < TimeSpan >]
```

Usage

```
Set-TransportConfig -MaxRetriesForLocalSiteShadow 3
```


-MaxRetriesForRemoteSiteShadow 4
-RejectMessageOnShadowFailure \$false
-SafetyNetHoldTime "3.00:00:00"
-ShadowHeartbeatFrequency "00:05:00"
-ShadowResubmitTimeSpan "02:00:00"
-ShadowMessageAutoDiscardInterval "3.00:00:00"

When working with shadow redundancy, Safety Net, and queues, you also want to consider:

- **ConnectionInactivityTimeout** Configured for each Send and Receive connector by using Set-SendConnector and Set-ReceiveConnector. Sets the maximum time that an open SMTP connection between servers can remain idle before timing out. This value must be smaller than the ConnectionTimeout value. For Send connectors, the default is 10 minutes. For Receive connectors, the default is five minutes for both the Transport service and the Front End Transport service on Mailbox servers, but only one minute for Edge Transport servers.
- **ConnectionTimeout** Configured for each Receive connector using Set-ReceiveConnector. Sets the maximum time that an SMTP connection can be open between servers, even if the source server is transmitting data. The default is 10 minutes for both the Transport service and the Front End Transport service on Mailbox servers, but only five minutes for Edge Transport servers.
- **MessageExpirationTimeout** Configured for the Transport service on each Mailbox server using Set-TransportService. Specifies how long a message can remain in a queue before it expires. The default value is two days.

When configuring these settings, you'll want to consider the relative latency and speed of the network as well as level of messaging traffic. If a slow or congested network has high latency, you may need to configure higher timeout values. Keep in mind, however, that each open connection uses resources and that each connector allows a finite number of open connections. By default, with Send connectors, the maximum number of connections is 1,000 and the maximum number of connections per domain is 20.

Configuring Safety Net

All Mailbox servers use Safety Net to maintain a queue of messages that were recently delivered to recipients. As discussed in "Exchange Server Message Queues" in Chapter 17, "Managing Exchange Organizations," and in the previous section of this chapter, this feature works in conjunction with shadow redundancy. The primary server that sends a message maintains the primary Safety Net queue while a second server, called the shadow server, maintains the shadow Safety Net queue.

In Exchange Management Shell, you configure Safety Net with these parameters in mind:

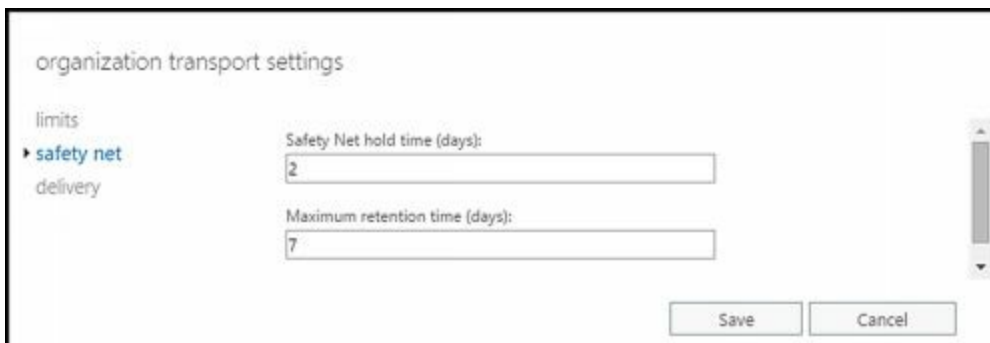
- **SafetyNetHoldTime** An organization-wide option configured for Set-TransportConfig. Specifies how long a successfully processed message is retained in the Safety Net queue. The default value is two days. Unacknowledged shadow messages expire after the sum of the SafetyNetHoldTime and the

MessageExpirationTimeout elapses. Set this value according to the size of your Exchange implementation and the relative amount of latency on the network. For example, a large global organization might want to set a longer time span, although the default may suffice for a smaller organization.

- **ReplayLagTime** Configured on individual mailbox database copies for Set-MailboxDatabaseCopy. Specifies how long the Exchange Replication service waits before replaying log files that have been copied to the passive database copy. By default, this option is not set. To ensure no data is lost and messages are available for resubmittal from the Safety Net queue, the replay lag time must be less than or equal to the safety net hold time.
- **MessageExpirationTimeout** Configured for the Transport service on each Mailbox server using Set-TransportService. Specifies how long a message can remain in a queue before it expires. The default value is two days.
- **ShadowRedundancyEnabled** Set using the Set-TransportConfig cmdlet. Enables or disables shadow redundancy for the Exchange organization. As Safety Net relies on shadow redundancy, you also disable Safety Net if you disable shadow redundancy.

You can use Exchange Admin Center to view or change the Safety Net hold time as well:

1. Select Mail Flow in the Navigation menu, and then select either Receive Connectors or Send Connectors.
2. In the main pane, select the More button (**⋮**), and then click Organization Transport Settings. This displays the Organization Transport Settings dialog box.
3. Click Safety Net. Finally, in the Safety Net Hold Time text box, enter the number of days that messages should be held in Safety Net queues and then click Save.



Enabling Anti-Spam Features

By default, Edge Transport servers have anti-spam features enabled and Mailbox servers do not. In an Exchange organization with Edge Transport servers, this is the desired configuration: you want your Edge Transport servers to run anti-spam filters on messages before they are routed into the Exchange organization. After Edge Transport servers have filtered messages, you don't need to filter them again, which is why Mailbox servers have this feature disabled.

If your organization doesn't use Edge Transport servers and has only Mailbox servers, you can enable the anti-spam features on Mailbox servers that receive messages from the Internet so that you can filter incoming messages for spam. However, if incoming

mail has any prior anti-spam filtering, you don't need to filter messages again.

The following anti-spam agents are available for the Transport service on Mailbox servers to use:

- [Content Filter agent](#)
- [Protocol Analysis agent](#)
- [Recipient Filter agent](#)
- [Sender Filter agent](#)
- [Sender ID agent](#)

You can install and configure these agents by doing the following:

1. Log on to the Mailbox server you want to configure.
2. In Exchange Management Shell, run the following command:

```
& $env:ExchangeInstallPath\Scripts\Install-AntiSpamAgents.ps1
```

3. After you install the anti-spam agents, you must restart the Exchange Transport service. In the shell, you can do this by running the following command:

```
Restart-service MExchangeTransport
```

4. Repeat Steps 1 - 3 for each Mailbox server that should filter messages.
5. Configure organization-wide transport settings that identify any internal SMTP servers that should be ignored by the Sender ID agent. Typically, this includes any Mailbox server in which you've enabled the anti-spam features. Use the `-InternalSMTPServers` parameter of `Set-TransportConfig` to identify each server by its IPv4 address. Here are examples:

```
Set-Transportconfig -InternalSMTPServers @{Add="192.168.10.52"}
```

```
Set-Transportconfig -InternalSMTPServers  
@{Add="192.168.10.52","192.168.10.64"}
```

6. You can verify that the servers were added by running the following command:

```
Get-TransportConfig | fl InternalSMTPServers
```

Once you've installed the anti-spam agents, you can enable or disable the anti-spam features on Mailbox servers by using the `Set-TransportService` cmdlet. To enable these features, set the `-AntispamAgentsEnabled` parameter to `$true`. To disable these features, set the `-AntispamAgentsEnabled` parameter to `$false`.

The following example shows how you can enable anti-spam features on a Mailbox server named `CorpSvr127`:

```
Set-TransportService -Identity 'CorpSvr127' -AntispamAgentsEnabled $true
```

Next you need to restart the Microsoft Exchange Transport service on the server. In the shell, you can do this by running the following command:

```
Restart-service MExchangeTransport
```

You can now configure the transport server's anti-spam features as discussed in the "Enabling Anti-Spam Features" section in Chapter 21. When you turn on anti-spam features, a transport server can automatically get updates for spam signatures, IP reputation, and anti-spam definitions through automatic updates, as long as you've done the following:

- [Conformed to Microsoft's licensing requirements](#)
- [Enabled Automatic Updates for use on the server](#)
- [Specifically enabled and configured anti-spam updates](#)

To obtain anti-spam updates through automatic updates, Microsoft requires an Exchange Enterprise Client Access License (CAL) for each mailbox user. You can configure automatic updates by using the Windows Update utility in Control Panel. Press Windows key + I, click Control Panel\Security, and then click Windows Update to start this utility. You can also configure Automatic Updates through Group Policy.

Subscribing Edge Transport Servers

When your Exchange organization uses Edge Transport servers and you want to use the Edge Synchronization feature, you must subscribe the Edge Transport server to your Exchange organization prior to performing other configuration tasks on the Edge Transport server. Creating a subscription allows the Microsoft Exchange EdgeSync service running on designated Mailbox servers to establish one-way replication of recipient and configuration information from your internal Active Directory database to the AD LDS database on an Edge Transport server. After you create an Edge subscription, synchronization is automatic. If problems occur, however, you can force synchronization or remove the Edge subscription.

Creating an Edge Subscription

A subscribed Edge Transport server receives the following from the EdgeSync service:

- [Send connector configurations](#)
- [Accepted domain configurations](#)
- [Remote domain configurations](#)
- [Safe Senders lists](#)
- [Recipients](#)

Any manually configured accepted domains, message classifications, remote domains, and Send connectors are deleted as part of the subscription process, and the related Exchange management interfaces are locked out as well. To manage these features after a subscription is created, you must do so within the Exchange organization and have the EdgeSync service update the Edge Transport server.

Also as part of the subscription process, you must select an Active Directory site for the subscription. The Mailbox server or servers in the site are the servers responsible for replicating Active Directory information to the Edge Transport server.

You can create a subscription for an Edge Transport server by completing the following steps:

1. Log on to the Edge Transport server for which you are creating a subscription by using an administrator account.
2. At Exchange Management Shell prompt, run the following command:

New-EdgeSubscription -filename "C:\EdgeSubscriptionExport.xml"

3. When prompted, confirm that it's okay to delete any manually configured accepted domains, message classifications, remote domains, and Send connectors by pressing A (which answers Yes to all deletion prompts).
4. Copy the EdgeSubscriptionExport.xml file to a Mailbox server in your Exchange organization.
5. Log on to a Mailbox server in your Exchange organization by using an account

with Exchange administration privileges.

6. On the Mailbox server, import the Edge Subscription file by running the following command:

New-EdgeSubscription -filename *FilePath*

Where *FilePath* specifies the full file path to the Edge Subscription file, such as:

New-EdgeSubscription -filename "C:\EdgeSubscriptionExport.xml"

Initial synchronization will begin, as discussed in “Synchronizing Edge Subscriptions.” Note that Mailbox servers in the Active Directory site must be able to resolve the IP addresses for the Edge Transport server. You need to ensure that subnets have been created in Active Directory Sites And Services and that DNS is configured to resolve the fully qualified domain name of the Edge Transport server. Mailbox servers in the site must also be able to connect to the Edge Transport server over TCP port 50636.

Listing 21-3 provides the syntax and usage for using the `New-EdgeSubscription` cmdlet to start a subscription. By default, the `-CreateInboundSendConnector` parameter is set to `$true`, which ensures that a Send connector from the Edge Transport server to Mailbox servers is created. By default, the `-CreateInternetSendConnector` parameter is set to `true`, which ensures that a Send connector to the Internet is created.

LISTING 21-3 New-EdgeSubscription cmdlet syntax and usage

Syntax

```
New-EdgeSubscription -FileName FilePath
-Site SiteName [-AccountExpiryDuration < TimeSpan >]
[-CreateInboundSendConnector <$true | $false>]
[-CreateInternetSendConnector <$true | $false>]
[-DomainController DCName ] [-FileData ByteStr ] [-Force <$true | $false>]
```

Usage

```
New-EdgeSubscription -FileName "Z:\EdgeSubscriptionExport.xml"
-Site "Default-First-Site-Name"
-CreatInboundSendConnector $true
-CreatInternetSendConnector $true
```

Getting Edge Subscription Details

You get information about Edge subscriptions using the `Get-EdgeSubscription` cmdlet. Each Edge subscription is listed by Edge Transport server name and associated Active Directory site. If you do not provide an identity with this cmdlet, configuration information for all Edge Subscriptions is returned.

LISTING 21-4 Get-EdgeSubscription cmdlet syntax and usage

Syntax

```
Get-EdgeSubscription -Identity EdgeTransportServerName
[-DomainController DCName ]
```

Usage

```
Get-EdgeSubscription -Identity "EdgeSvr04"  
Get-EdgeSubscription -Identity "EdgeSvr04" | fl
```

Synchronizing Edge Subscriptions

During the configuration of an Edge subscription, you specified an Active Directory site to associate with the subscription. Mailbox servers in this site run the EdgeSync service and are responsible for synchronizing configuration data between Active Directory Domain Services and AD LDS on the Edge Transport server. By default, the EdgeSync service synchronizes configuration data hourly and recipient data every four hours.

If you've just created a new subscription and synchronization has occurred, you should verify that replication is taking place as expected by completing the following steps:

1. On the Edge Transport server, start Exchange Management Shell.
2. Verify that a Send connector was created to send Internet mail by typing the command **get-sendconnector**. As shown in the following example and sample output, you should see an Inbound connector and an Internet connector for EdgeSync:

```
get-sendconnector
```

Identity	AddressSpaces	Enabled
Primary Send Connector	{SMTP:*.tvpress.com;1}	True
SD1 Send Connector	{SMTP:*.imaginedlands.com;1}	True
EdgeSync - Seattle-First-Site to Int	{smtp:*;100}	True
EdgeSync - Inbound to Seattle-First-	{smtp:--;100}	True

3. Verify that there is at least one entry for accepted domains by typing **get-accepteddomain** as shown in the following example and sample output:

```
get-accepteddomain
```

Name	DomainName	DomainType	Default
tvpress.com	tvpress.com	Authoritative	True

If you suspect there is a problem with synchronization and you want to start immediate synchronization of configuration data for all Edge subscriptions, complete the following steps:

4. Start Exchange Management Shell.
5. At the prompt, type the following command

```
start-edgesynchronization -Server ServerName
```

where **ServerName** is the name of the Mailbox server on which you want to run the command, such as:

```
start-edgesynchronization -Server mailserv25
```


If you are running the command on the Mailbox server, you can omit the `-Server` parameter.

Verifying Edge Subscriptions

The easiest way to verify the subscription status of Edge Transport servers is to run the `Test-EdgeSynchronization` cmdlet. This cmdlet provides a report of the synchronization status, and you also can use it to verify that a specific recipient has been synchronized to the Active Directory Lightweight Directory Service on an Edge Transport server.

Listing 21-5 provides the syntax and usage for the `Test-EdgeSynchronization` cmdlet. By default, the cmdlet verifies configuration objects and recipient objects. To have the cmdlet verify only configuration data, set `-ExcludeRecipientTest` to `$true`. Use the `-VerifyRecipient` parameter to specify the email address of a recipient to verify.

LISTING 21-5 `Test-EdgeSynchronization` cmdlet syntax and usage

Syntax

```
Test-EdgeSynchronization [-ExcludeRecipientTest <$true | $false>]
[-DomainController DCName ] [-FullCompareMode <$true | $false>]
[-MaxReportSize < MaxNumberOfObjectsToCheck | Unlimited>]
[-MonitoringContext <$true | $false>] [-TargetServer EdgeServer ]
```

```
Test-EdgeSynchronization -VerifyRecipient RecipientEmailAddress
[-DomainController DCName ]
```

Usage

```
Test-EdgeSynchronization -ExcludeRecipientTest
```

```
Test-EdgeSynchronization -MaxReportSize 500
```

```
Test-EdgeSynchronization -VerifyRecipient "williams@tvpress.com"
```

```
Test-EdgeSynchronization -TargetServer CorpServer73.tvpress.com
```

Example and sample output

```
test-edgesynchronization
```

```
RunspaceId           :
UtcNow                : 3/18/2016 3:11:22 PM
Name                  : CORPSEVER73
LeaseHolder           : CN=MAILSERVER25,CN=Servers,CN=Exchange
Administrative Group (FYDIBOHF23SPDLT),CN=Administrative Groups,
CN=First Organization,CN=Microsoft Exchange,CN=Services,CN=Configuration,
DC=cpan1,DC=com
LeaseType             : Option
ConnectionResult      : Succeeded
FailureDetail         :
LeaseExpiryUtc        : 3/18/2016 3:06:30 PM
```

LastSynchronizedUtc : 3/18/2016 3:11:59 PM
CredentialStatus : Synchronized
TransportServerStatus : Synchronized
TransportConfigStatus : Synchronized
AcceptedDomainStatus : Synchronized
RemoteDomainStatus : Synchronized
SendConnectorStatus : Synchronized
MessageClassificationStatus : Synchronized
RecipientStatus : Synchronized
CredentialRecords : Number of credentials 85
CookieRecords : Number of cookies 27

Removing Edge Subscriptions

If you replace or decommission an Edge Transport server, you no longer need the related Edge subscription and can remove it. Removing an Edge subscription

- Stops synchronization of information from the Active Directory Domain Service to AD LDS.
- Removes all the accounts that are stored in AD LDS.
- Removes the Edge Transport server from the source server list of any Send connector.

In Exchange Management Shell, you can remove an Edge Subscription by passing the identity of the subscription to remove to the Remove-EdgeSubscription cmdlet. Listing 21-6 provides the syntax and usage.

LISTING 21-6 Remove-EdgeSubscription cmdlet syntax and usage

Syntax

```
Remove-EdgeSubscription -Identity EdgeTransportServerName  
[-DomainController DCName ] [-Force <$true | $false>]
```

Usage

```
Remove-EdgeSubscription -Identity "EdgeSvr04"
```

Chapter 22. Maintaining Mail Flow

With Exchange 2016, mail flow occurs through a collection of services, connections, components, and queues that work together as part of the transport pipeline. The Front End Transport service acts as a stateless proxy for all inbound and outbound external SMTP traffic. The Microsoft Exchange Transport service (running as a back-end component) categorizes messages, inspects their content, and queues them for delivery or submission.

Message delivery is handled by the Mailbox Transport Delivery service, and message submission is handled by the Mailbox Transport Submission service—both of which are components of the Microsoft Exchange Transport service. Although the transport pipeline is critical to mail flow, many other factors also affect mail flow in an Exchange organization, including the configuration of message processing speeds, message throttling, accepted domains, email address policies, journal rules, remote domains, filters, and transport rules.

Managing Message Routing and Delivery

To support message routing and delivery, Mailbox and Edge Transport servers maintain a few special directories:

- **Pickup** A folder to which users and applications can manually create and submit new messages for delivery.
- **Replay** A folder for messages bound for or received from non-SMTP mail connectors.

The sections that follow discuss how the Pickup and Replay directories are used and configured and also look at the related concepts of message throttling and back pressure.

Understanding Message Pickup and Replay

When a Mailbox or an Edge Transport server receives incoming mail from a server using a non-SMTP connector, it stores the message in the Replay directory and then resubmits it for delivery using SMTP. When a Mailbox or an Edge Transport server has a message to deliver to a non-SMTP connector, it stores the message in the Replay directory and then resubmits it for delivery to the foreign connector. In this way, messages received from non-SMTP connectors are processed and routed, and messages to non-SMTP connectors are delivered.

Your Transport servers automatically process any correctly formatted .eml message file copied into the Pickup directory. Exchange considers a message file that is copied into the Pickup directory to be correctly formatted if it meets the following conditions:

- Is a text file that complies with the basic SMTP message format and can also use Multipurpose Internet Mail Extension (MIME) header fields and content
- Has an .eml file name extension, zero or one email address in the Sender field, and one or more email addresses in the From field
- Has at least one email address in the To, Cc, or Bcc fields and a blank line between the header fields and the message body

Transport servers check the Pickup directory for new message files every five seconds. Although you can't modify this polling interval, you can adjust the rate of message file processing by using the `-PickupDirectoryMaxMessagesPerMinute` parameter on the `Set-TransportServer` cmdlet. The default value is 100 messages per minute. When a transport server picks up a message, it checks the message against the maximum message size, the maximum header size, the maximum number of recipients, and other messaging limits.

For the Pickup directory, the maximum message size is 10 megabytes (MB), the maximum header size is 64 kilobytes (KB), and the maximum number of recipients is 100 by default. As may be necessary to meet the needs of your organization, you can change these limits by using the `Set-TransportServer` cmdlet. If a message file doesn't exceed any assigned limits, the Transport server renames the message file by using a

.tmp extension, and then converts the .tmp file to an email message. After the message is successfully queued for delivery, the Transport server issues a “close” command and deletes the .tmp file from the Pickup directory.

REAL WORLD Header fields are plain text, and each character of text is 1 byte. The size of the header is determined by the total number of header fields and characters in each header field. Organization X-headers, forest X-headers, and routing headers are removed from messages in the Pickup directory. On the other hand, routing headers are preserved in the Replay directory, and organization X-headers and forest X-headers also are preserved if an X-CreatedBy header field indicates the headers were created by Exchange 2016 (meaning the field value is set to MExchange15).

Your Transport servers automatically process any correctly formatted .eml message file copied into the Replay directory. Exchange considers a message file that is copied into the Replay directory to be correctly formatted if it meets the following conditions:

- Is a text file that complies with the basic SMTP message format and can also use MIME header fields and content.
- Has an .eml file name extension, and its X-header fields occur before all regular header fields.
- Has a blank line between the header fields and the message body.

Transport servers check the Replay directory for new message files every five seconds. Although you can't modify this polling interval, you can adjust the rate of message file processing by using the `-PickupDirectoryMaxMessagesPerMinute` parameter of the `Set-TransportServer` cmdlet. This parameter controls the rate of processing for both the Pickup directory and the Replay directory. The Transport server renames the message file by using a .tmp extension, and then converts the .tmp file to an email message. After the message is successfully queued for delivery, the server issues a “close” command and deletes the .tmp file from the Replay directory.

Exchange considers any improperly formatted email messages received in the Pickup or Replay directory to be undeliverable and renames them from the standard message name (*MessageName*.eml) to a bad message name (*MessageName*.bad). Because this is considered a type of message-processing failure, a related error is also generated in the event logs. In addition, if you restart the Microsoft Exchange Transport service when .tmp files are in the Pickup directory, Replay directory, or both directories, all .tmp files are renamed as .eml files and are reprocessed, which can lead to duplicate message transmissions.

Configuring and Moving the Pickup and Replay Directories

Because of the way message pickup and replay works, Transport servers do not perform security checks on messages submitted through these directories. This means that if you've configured anti-spam, antivirus, sender filtering, or recipient filtering actions on

a Send connector, those checks are not performed on the Pickup or Replay directory. To ensure that the Pickup and Replay directories are not compromised by malicious users, specific security permissions that must be tightly controlled are applied.

For the Pickup and Replay directories, you must configure the following permissions:

- Full Control for Administrator
- Full Control for Local System @techjob
- Read, Write, and Delete Subfolders and Files for Network Service

When you have a need to balance the load across a server's disk drives or ensure ample free space for messages, you can move the Pickup and Replay directories to new locations. By using Set-TransportServer with the `-PickupDirectoryPath` parameter, you can move the location of the Pickup directory. Move the location of the Replay directory by using the `-ReplayDirectoryPath` parameter on the Set-TransportServer cmdlet. With either parameter, successfully changing the directory location depends on the rights that are granted to the Network Service account on the new directory location and whether the new directory already exists. Keep the following in mind:

- If the new directory does not already exist and the Network Service account has the rights to create folders and apply permissions at the new location, the new directory is created and the correct permissions are applied to it.
- If the new directory already exists, the existing folder permissions are not checked or changed. Exchange assumes you've already set the appropriate permissions.

Listing 22-1 provides the syntax and usage for moving the Pickup and Replay directories. If you want to move both the Pickup and Replay directories, you should do so in two separate commands.

LISTING 22-1 Changing the Pickup directory

Syntax

```
Set-TransportServer -Identity ServerIdentity  
[-PickupDirectoryPath LocalFolderPath]  
[-ReplayDirectoryPath LocalFolderPath]
```

Usage

```
Set-TransportServer -Identity "CorpSvr127"  
-PickupDirectoryPath "g:\Pickup"
```

Changing the Message Processing Speed

By default, Transport servers simultaneously and separately process the Pickup and Replay directories. Transport servers scan the Pickup and Replay directories for new message files once every five seconds (or 12 times per minute), and they process messages copied to either directory at a rate of 100 messages per minute, per directory. Because the polling interval is not configurable, the maximum number of messages that can be processed in either the Pickup or Replay directory during each polling interval, by default, is approximately eight (100 messages per minute divided by 12 messages

processed per minute).

Although the polling interval is not configurable, the maximum number of messages that can be processed during each polling interval is configurable. You assign the desired processing rate by using the `-PickupDirectoryMaxMessagesPerMinute` parameter, because this processing speed is used with both the Pickup and Replay directories. You might want to adjust the message processing rate in these situations:

- If the server is unable to keep up with message processing, you might want to decrease the number of messages processed per minute to reduce processor and memory utilization.
- If the server is handling message transport for a large organization and you are seeing delays in message transport because of an abundance of messages in the Pickup directory, Replay directory, or both directories, you might want to increase the number of messages processed per minute, as long as the server can handle the additional workload.

You assign the desired processing rate by using the `-PickupDirectoryMaxMessagesPerMinute` parameter of the `Set-TransportServer` cmdlet, as shown in Listing 22-2, and this processing speed is used with both the Pickup and Replay directories. Your Transport server then attempts to process messages in each directory independently at the rate specified. You can use a per-minute message processing value between 1 and 20,000.

LISTING 22-2 Changing the message processing speed

Syntax

```
Set-TransportServer -Identity ServerIdentity  
[-PickupDirectoryMaxMessagesPerMinute Speed]
```

Usage

```
Set-TransportServer -Identity "CorpSvr127"  
-PickupDirectoryMaxMessagesPerMinute "500"
```

Configuring Messaging Limits for the Pickup Directory

The Pickup directory is used by administrators to test mail flow and by applications that create and submit their own messages. If applications are generating messages with expanded headers, such as when there are many recipients in `To:`, `Cc:`, and `Bcc:` header fields, you may need to modify the messaging limits for the Pickup directory.

You can set messaging limits for the Pickup directory for message header sizes and maximum recipients per message. The default message header size is 64 KB, which allows for 65,536 characters in the header. To change this setting, you can set the `-PickupDirectoryMaxHeaderSize` parameter of the `Set-TransportServer` cmdlet to the desired size. The valid input range for this parameter is 32,768 to 2,147,483,647 bytes. When you specify a value, you must qualify the units for that value by ending with one of the following suffixes:

- B for bytes
- KB for kilobytes
- MB for megabytes
- GB for gigabytes

The following example sets the maximum header size to 256 KB:

```
Set-TransportServer –Identity MailServer48  
–PickupDirectoryMaxHeaderSize "256KB"
```

The default maximum recipients per message is 100. To change this setting, you can set the `–PickupDirectoryMaxRecipientsPerMessage` parameter of the `Set-TransportServer` cmdlet to the desired size. The valid input range for this parameter is 1 to 10,000. The following example sets the maximum recipients to 500:

```
Set-TransportServer –Identity MailServer48  
–PickupDirectoryMaxRecipientsPerMessage "500"
```

Configuring Message Throttling

Message throttling sets limits on the number of messages and connections that can be processed by a Mailbox or an Edge Transport server. These limits are designed to prevent the accidental or intentional inundation of transport servers and help ensure that transport servers can process messages and connections in an orderly and timely manner. Throttling works in conjunction with size limits on messages that apply to header sizes, attachment sizes, and number of recipients. Although the default throttling settings work in a typical messaging environment, you may need to modify these settings as your organization grows, especially if users or applications create and send a lot of email messages.

On Mailbox and Edge Transport servers, you can set some message throttling options in the Exchange Admin Center by using the options on the Transport Limits page in the transport server's Properties dialog box. In Exchange Management Shell, you can configure all message throttling options by using `Set-TransportServer` and related parameters.

- **MaxConcurrentMailboxDeliveries** Sets the maximum number of delivery threads that can be open at the same time to deliver messages to mailboxes. The default value is 20.
- **MaxConcurrentMailboxSubmissions** Sets the maximum number of delivery threads that can be open at the same time to accept messages from mailboxes. The default value is 20.
- **MaxConnectionRatePerMinute** Sets the maximum rate at which new inbound connections can be opened to any Receive connectors that exist on the server. The default value is 1,200 connections per minute.
- **MaxOutboundConnections** Sets the maximum number of concurrent outbound connections that can be open at the same time for Send connectors. The default value is 1,000.

- **MaxPerDomainOutboundConnections** Sets the maximum number of connections that can be open to any single remote domain for any available Send connectors. The default value is 20.

With Set-SendConnector, you can configure throttling by using ConnectionInactivityTimeout. This parameter sets the maximum idle time before an open SMTP connection is closed. The default value is 10 minutes.

With Set-ReceiveConnector, you can configure throttling by using the following parameters:

- **ConnectionInactivityTimeout** Sets the maximum idle time before an open SMTP connection is closed. The default value is five minutes for a Mailbox and one minute for an Edge Transport.
- **ConnectionTimeout** Sets the maximum time that an SMTP connection can remain open, even if it is active. The default value is 10 minutes for a Mailbox and five minutes for an Edge Transport. ConnectionTimeout must be longer than ConnectionInactivityTimeout.
- **MaxInboundConnection** Sets the maximum number of simultaneous inbound SMTP connections. The default value is 5,000.
- **MaxInboundConnectionPercentagePerSource** Sets the maximum number of simultaneous inbound SMTP connections from a single source server. The value is expressed as the percentage of available remaining connections on a Receive connector (as defined by the –MaxInboundConnection parameter). The default value is 2 percent.
- **MaxInboundConnectionPerSource** Sets the maximum number of simultaneous inbound SMTP connections from a single source messaging server. The default value is 100.
- **MaxProtocolErrors** Sets the maximum number of SMTP protocol errors allowed before a Receive connector closes a connection with a source messaging server. The default value is five.
- **TarpitInterval** Sets artificial delay in SMTP responses in cases in which unwelcome messages are being received from anonymous connections. The default value is five seconds.

Understanding Back Pressure

Back pressure limits overutilization of system resources on a Mailbox or an Edge Transport server. Transport servers monitor key system resources to determine usage levels. If usage levels exceed a specified limit, the server stops accepting new connections and messages to prevent server resources from being completely overwhelmed and to enable the server to deliver the existing messages. When usage of system resources returns to a normal level, the server accepts new connections and messages. Resources monitored as part of the back pressure feature include:

- Free space on hard disk drives that store the message queue database transaction logs.
- Free space on the hard disk drives that store the message queue database.

- The amount of memory used by all processes.
- The amount of memory used by the Edgetransport.exe process.
- The number of uncommitted message queue database transactions that exist in memory.

Levels of usage are defined as normal, medium, or high. With the normal level, the resource is not overused, and the server accepts new connections and messages. With the medium level, the resource is slightly overused, and limited back pressure is applied, allowing mail from senders in the authoritative domain to continue being sent while the server rejects new connections and messages from other sources. With the high level, the resource is severely overused and full back pressure is applied, meaning message flow stops and the server rejects all new connections and messages.

You have limited control over how back pressure is applied. Some related settings can be configured in the Edgetransport.exe.config file on Edge Transport servers; however Microsoft recommends that you don't change the default settings.

Creating and Managing Accepted Domains

Any SMTP namespace for which an Exchange organization sends or receives email is an accepted domain. Accepted domains include domains for which the Exchange organization is authoritative, as well as domains for which the Exchange organization relays mail.

Understanding SMTP Domains

An Exchange organization can have more than one SMTP domain, and the set of email domains your organization uses are its authoritative domains. An accepted domain is considered authoritative when the Exchange organization hosts mailboxes for recipients in this SMTP domain. Transport servers should always accept email that is addressed to any of the organization's authoritative domains. By default, when you install the first Mailbox server, one accepted domain is configured as authoritative for the Exchange organization, and this default accepted domain is based on the FQDN of your forest root domain.

Often an organization's internal domain name might differ from its external domain name. You must create an accepted domain to match your external domain name. You must also create an email address policy that assigns your external domain name to user email addresses. For example, your internal domain name might be tvpress.local, whereas your external domain name is tvpress.com. When you configure DNS, the DNS MX records for your organization will reference tvpress.com, and you will want to assign this SMTP namespace to users by creating an email address policy.

When email is received from the Internet by a Transport server and the recipient of the message is not a part of your organization's authoritative domains, the sending server is trying to relay messages through your Transport servers. To prevent abuse of your servers, Transport servers reject all email that is not addressed to a recipient in your organization's authoritative domains. However, at times you might need to relay email messages from another domain, such as messages from a partner or subsidiary, in which case, you can configure accepted domains as relay domains. When your Transport servers receive the email messages for a configured relay domain, they will relay the messages to an email server in that domain.

Two options are available for configuring relay domains:

- [Internal relay domain](#)
- [External relay domain](#)

You configure an internal relay domain when there are contacts from the relay domain in the global address list. If your organization contains more than one forest and has configured global address list synchronization, the SMTP domain for one forest can be configured as an internal relay domain in a second forest. Messages from the Internet that are addressed to recipients in internal relay domains are received and processed by

your Edge Transport servers. These messages are then relayed to your Mailbox servers, which, in turn, route the messages to the Mailbox servers in the recipient forest. Configuring an SMTP domain as an internal relay domain ensures that all email addressed to the relay domain are accepted by your Exchange organization.

You configure an external relay domain when you want to relay messages to an email server that is both outside your Exchange organization and outside the boundaries of your organization's network perimeter. For this configuration to work, your DNS servers must have an MX record for the external relay domain that references a public IP address for the relaying Exchange organization. When your Edge Transport servers receive the messages for recipients in the external relay domain, they route the messages to the mail server for the external relay domain. You must also configure a Send connector from the Edge Transport server to the external relay domain. The external relay domain can also be using your organization's Edge Transport server as a smart host for outgoing mail.

You also can configure accepted domains for Microsoft Exchange Online. In this case, accepted domains can either be authoritative or internal relay domains. Although you manage previously defined domains in Exchange Admin Center under Mail Flow>Accepted Domains, you must initially define domains in Office 365 Admin Center by using the Domains > Add A Domain option.

If you are working in a hybrid organization, you'll find that the Hybrid Configuration Wizard adds an accepted domain to the on-premises organization to enable hybrid mail flow. This domain, called the coexistence domain, is added as a secondary proxy domain to any email address policies that have primary SMTP address templates for domains selected in the wizard. By default, the coexistence domain is *YourDomain*.mail.onmicrosoft.com.

Viewing Accepted Domains

To view the accepted domains configured for your organization, complete the following steps:

1. In the Exchange Admin Center, select Mail Flow in the Navigation menu, and then select Accepted Domains.
2. In the main pane, accepted domains are listed by name, SMTP domain name, and domain type. The domain type is listed as Authoritative, External Relay, or Internal Relay as shown in Figure 22-1.

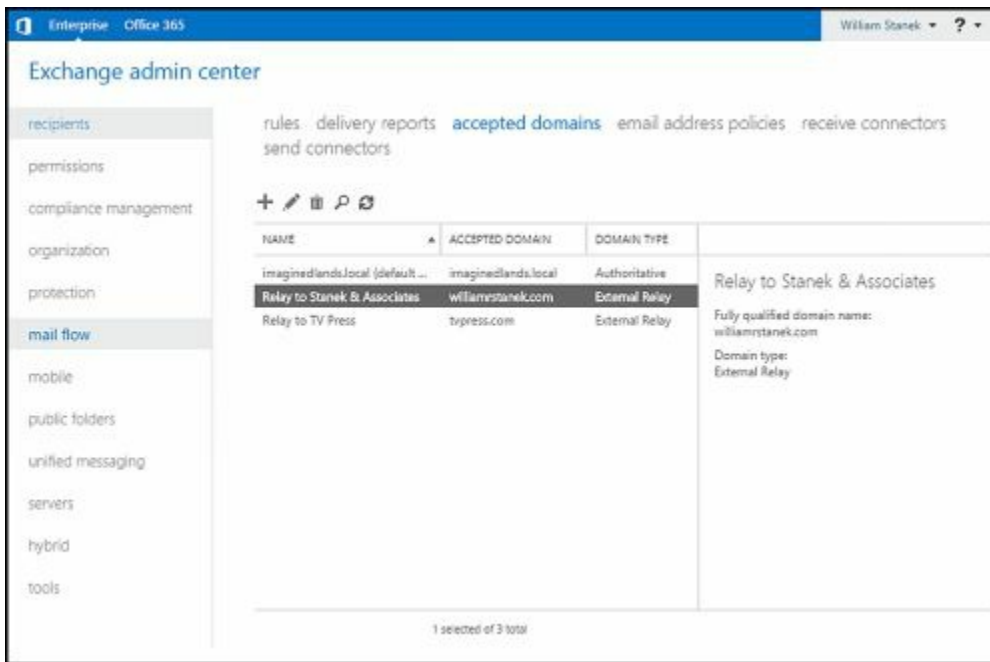


FIGURE 22-1 View accepted domains.

You can use the `Get-AcceptedDomain` cmdlet to list accepted domains or to get information on a particular accepted domain as well. If you do not provide an identity with this cmdlet, configuration information for all accepted domains is displayed. Listing 22-3 provides the syntax and usage, as well as sample output, for the `Get-AcceptedDomain` cmdlet.

LISTING 22-3 `Get-AcceptedDomain` cmdlet syntax and usage

Syntax

```
Get-AcceptedDomain [-Identity DomainIdentity]
[-DomainController DCName ] [-Organization OrganizationId ]
```

Usage

```
Get-AcceptedDomain
Get-AcceptedDomain -Identity "imaginedlands.local"
Get-AcceptedDomain | Where {$_.DomainType -eq 'Authoritative'}
```

Example Output

Name	DomainName	DomainType	Default
imaginedlands.local	imaginedlands.local	Authoritative	True
Relay to TV Press	tvpress.com	ExternalRelay	False
Relay to Stanek & Associates	williamrstanek.com	ExternalRelay	False

Creating Accepted Domains

To create accepted domains for your organization, complete the following steps:

1. In the Exchange Admin Center, select Mail Flow in the Navigation menu, and

then select Accepted Domains. Next, click Add () to open the New Accepted Domain dialog box, as shown in Figure 22-2.

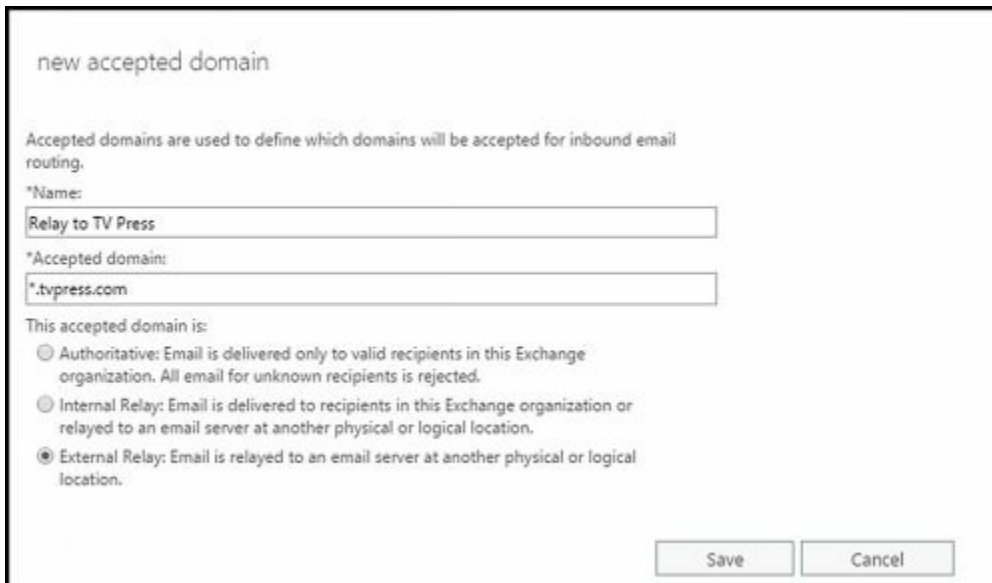


FIGURE 22-2 Create a new accepted domain.

2. Use the Name text box to identify the accepted domain. You can use a descriptive name that identifies the purpose of the accepted domain or simply enter the actual SMTP domain name.
3. In the Accepted Domain text box, type the SMTP domain name for which the Exchange organization will accept email messages. If you want to accept email for the specified domain only, enter the full domain name, such as **imaginedlands.com**. If you want to accept email for the specified domain and child domains, type * (a wildcard character), then a period, and then the domain name, such as ***.imaginedlands.com**.

NOTE Only domain names you specify can be used as part of an email address policy. Because of this, if you want to use a subdomain as part of an email address policy, you must either explicitly configure the subdomain as an accepted domain or use a wildcard character to include the parent domain and all related subdomains.

4. Select one of the following options to set the accepted domain type:
 - **Authoritative Domain** Email is delivered to a recipient in this exchange organization.
 - **Internal Relay Domain** Email is relayed to an email server in another Active Directory forest in the organization.
 - **External Relay Domain** Email is relayed to an email server outside the organization by the Edge Transport server.
5. Click Save to create the accepted domain.

In Exchange Management Shell, you can use the `New-AcceptedDomain` cmdlet to create accepted domains. Listing 22-4 provides the syntax and usage.

LISTING 22-4 New-AcceptedDomain cmdlet syntax and usage

Syntax


```
New-AcceptedDomain -Name Name  
-DomainName DomainName  
-DomainType <Authoritative|InternalRelay|ExternalRelay>  
[-Organization OrganizationId]
```

Usage;

```
new-AcceptedDomain -Name "Relay to TV Press"  
-DomainName "*.tvpress.com"  
-DomainType "ExternalRelay"
```

Changing The Accepted Domain Type and Identifier

You can change an accepted domain's type and identifier by completing the following steps:

1. In the Exchange Admin Center, select Mail Flow in the Navigation menu, and then select Accepted Domains. Next, select the accepted domain you want to change, and then select Edit (). Or simply double-click the accepted domain.
2. In the Properties dialog box, enter a new identifier, use the options provided to change the accepted domain type as necessary.
3. Select the Make This The Default Domain checkbox to make the currently selected domain the default for the Exchange organization. The default accepted domain is used in the external postmaster email address and in encapsulated non-SMTP email addresses.
4. Click Save.



Relay to Stanek & Associates

Accepted domains are used to define which domains will be accepted for inbound email routing.

*Name:
Relay to Stanek & Associates

Accepted domain:
williamrstanek.com

This accepted domain is:

- Authoritative: Email is delivered only to valid recipients in this Exchange organization. All email for unknown recipients is rejected.
- Internal Relay: Email is delivered to recipients in this Exchange organization or relayed to an email server at another physical or logical location.
- External Relay: Email is relayed to an email server at another physical or logical location.

Make this the default domain.

Save Cancel

In Exchange Management Shell, you can use the Set-AcceptedDomain cmdlet to modify

accepted domains. Listing 22-5 provides the syntax and usage. Use the `AddressBookEnabled` parameter to enable recipient filtering for this accepted domain. You should set this parameter to `$true` only if all the recipients in this accepted domain are replicated to the AD LDS database on the Edge Transport servers. For authoritative domains and internal relay domains, the default value is `$true`. For external relay domains, the default value is `$false`.

LISTING 22-5 Set-AcceptedDomain cmdlet syntax and usage

Syntax


```
Set-AcceptedDomain -Identity AcceptedDomainIdentity  
[-AddressBookEnabled <$true | $false>] [-DomainController DCName]  
[-DomainType <Authoritative|InternalRelay|ExternalRelay>]  
[-MakeDefault <$true | $false>] [-Name Name]
```

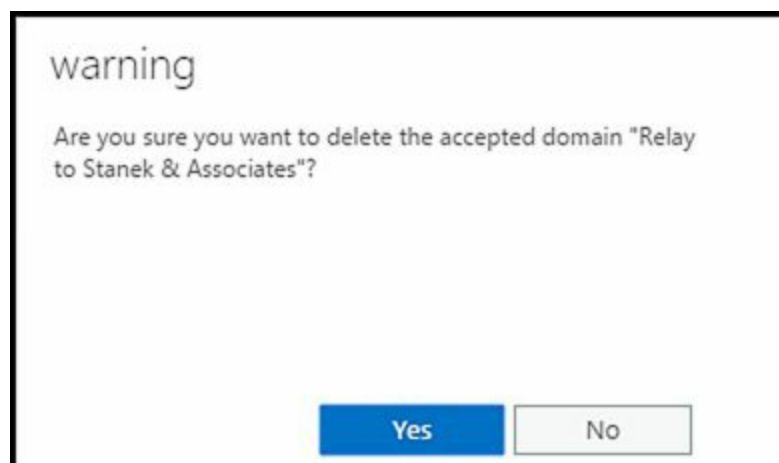
Usage

```
Set-AcceptedDomain -Identity "Relay to TV Press"  
-DomainType "ExternalRelay"
```

Removing Accepted Domains

You can remove an accepted domain that's no longer needed by completing the following steps:

1. In the Exchange Admin Center, select **Mail Flow** in the Navigation menu, and then select **Accepted Domains**. Next, select the accepted domain you want to delete, and then select **Delete** ().
2. When prompted to confirm, click **Yes**.



In Exchange Management Shell, you can use the `Remove-AcceptedDomain` cmdlet to remove accepted domains. Listing 22-6 provides the syntax and usage.

LISTING 22-6 Remove-AcceptedDomain cmdlet syntax and usage

Syntax

```
Remove-AcceptedDomain -Identity AcceptedDomainIdentity
```

[-DomainController **DCName**]

Usage

Remove-AcceptedDomain -Identity "Relay to TV Press"

Creating and Managing Remote Domains

In on-premises Exchange organizations, remote domain settings help you manage mail flow for most types of automated messages, including out-of-office messages, automatic replies, automatic forwarding, delivery reports, and nondelivery reports. Remote domain settings also control some automated message-formatting options, such as whether to display a sender's name on a message or only the sender's email address. Your Exchange organization has a default remote domain policy that sets the global defaults. You can create additional policies to create managed connections for specific remote domains as well.

Viewing Remote Domains

You can use the `Get-RemoteDomain` cmdlet to list remote domains or to get information on a particular remote domain. Remote domains are listed by name and the domain to which they apply. The Default remote domain applies to all remote domains, unless you override it with specific settings.

If you do not provide an identity with the `Get-RemoteDomain` cmdlet, configuration information for all remote domains is displayed. Listing 22-7 provides the syntax and usage, as well as sample output, for the `Get-RemoteDomain` cmdlet.

LISTING 22-7 `Get-RemoteDomain` cmdlet syntax and usage

Syntax

```
Get-RemoteDomain [-Identity DomainIdentity]  
[-DomainController DCName ] [-Organization OrgId ]
```

Usage

```
Get-RemoteDomain -Identity "imaginedlands.com"
```

Output

Name	DomainName	AllowedOOFTType
----	-----	-----
Default	*	External
The Magic Lands	*.imaginedlands.com	External

Creating Remote Domains

In Exchange Management Shell, you can use the `New-RemoteDomain` cmdlet to create remote domains. Use the `-Name` parameter to specify a descriptive name that identifies the purpose of the remote domain or simply enter the actual SMTP domain name.

Listing 22-8 provides the syntax and usage. The way you set the `-DomainName` parameter determines whether the remote domain includes subdomains. To manage connections for a specific domain, you simply provide the fully qualified name of the domain. You insert an asterisk and a period before the domain name to include the

domain and all child domains of the domain.

LISTING 22-8 New-RemoteDomain cmdlet syntax and usage

Syntax

```
New-RemoteDomain -Name Name -DomainName DomainName  
[-DomainController DCName] [-Organization OrgId]
```

Usage for parent domain only

```
New-RemoteDomain -Name "The Magic Lands Managed Connection"  
-DomainName "themagiclands.com"
```

Usage for parent domain and child domains

```
New-RemoteDomain -Name "The Magic Lands Managed Connection"  
-DomainName "*.themagiclands.com"
```

Configuring Messaging Options for Remote Domains

Remote domains are used to control how automated messages are used and to specify some types of messaging format options. In Exchange Management Shell, you can use the Set-RemoteDomain cmdlet to configure remote domains. Listing 22-9 provides the syntax and usage.

LISTING 22-9 Set-RemoteDomain cmdlet syntax and usage

Syntax

```
Set-RemoteDomain -Identity "RemoteDomainIdentity"  
[-AllowedOOFType <"External"|"InternalLegacy"|"ExternalLegacy"|"None">]  
[-AutoForwardEnabled <{$true | $false}>]  
[-AutoReplyEnabled <{$true | $false}>]  
[-CharacterSet "CharacterSet"]  
[-ContentType <"MimeHtmlText"|"MimeText"|"MimeHtml">]  
[-DeliveryReportEnabled <{$true | $false}>]  
[-DisplaySenderName <{$true | $false}>]  
[-DomainController DCName]  
[-LineWrapSize "Size"]  
[-MeetingForwardNotificationEnabled <{$true | $false}>]  
[-Name "Name"]  
[-NDREnabled <{$true | $false}>]  
[-NonMimeCharacterSet "CharacterSet"]  
[-TNEFEnabled <{$true | $false}>]
```

Usage

```
Set-RemoteDomain -Identity "The Magic Lands Managed Connection"  
-DeliveryReportEnabled $false
```

Use the -AllowedOOFType parameter to specify whether and how out-of-office messages are sent to the remote domain. The options are as follows:

- **None** Blocks all out-of-office messages.
- **External** Allows out-of-office messages to be received by the Exchange organization, but does not allow the organization's out-of-office messages to be sent.
- **ExternalLegacy** Allows out-of-office messages to be received by the Exchange organization and receipt of out-of-office messages generated by Microsoft Outlook 2003, Exchange 2003, or earlier.
- **InternalLegacy** Allows out-of-office messages to be sent from the Exchange organization and the sending of out-of-office messages generated by Outlook 2003, Exchange 2003, or earlier.

You also can specify how Exchange should format messages. Allow messaging options by setting the related parameters to \$true, or disallow messaging options by setting the related parameters to \$false. The options available are as follows:

- **-AutoReplyEnabled** Allows the sender to be notified that the message was received.
- **-AutoForwardEnabled** Allows Exchange Server to forward or deliver a duplicate message to a new recipient.
- **-DeliveryReportsEnabled** Allows Exchange Server to return delivery confirmation reports to the sender.
- **-MeetingForwardNotificationEnabled** Allows Exchange Server to forward or deliver a meeting notification to a new recipient.
- **-NDREnabled** Allows Exchange Server to return nondelivery confirmation reports to the sender.
- **-DisplaySenderName** Allows both the sender's name and email address to appear on outbound email messages.

By default, text word-wrapping is disabled, which means that Exchange does not enforce a maximum line length. If you'd like message text to wrap at a specific line length, you can set the **-LineWrapSize** parameter to the specific column position at which text wrapping should start, such as at 72 characters.

Use the **-ContentType** parameter to set the outbound message content type and formatting. The options are as follows:

- **MimeHTML** Converts messages to MIME messages with HTML formatting.
- **MimeText** Converts messages to MIME messages with text formatting.
- **MimeHtmlText** Converts messages to MIME messages with HTML formatting, except when the original message is a text message. Text messages are converted to MIME messages with text formatting.

If you want to send Transport Neutral Encapsulation Format (TNEF) message data to the remote domain rather than Exchange Rich Text Format, set **-TNEFEnabled** to \$true.

Removing Remote Domains

In Exchange Management Shell, you can use the **Remove-RemoteDomain** cmdlet to

remove remote domains. Listing 22-10 provides the syntax and usage.

LISTING 22-10 Remove-RemoteDomain cmdlet syntax and usage

Syntax

```
Remove-RemoteDomain -Identity RemoteDomainIdentity  
[-DomainController DCName]
```

Usage

```
Remove-RemoteDomain -Identity "The Magic Lands Managed Connection"
```


Chapter 23. Implementing Exchange Policies and Rules

Exchange provides email address policies, journal rules and transport rules as a means to affect mail flow by modifying the transport pipeline. Email address policies allow you to generate or rewrite email addresses automatically for each recipient in your organization based on specific criteria you set. Microsoft Exchange Server uses email address policies in two key ways:

- Whenever you create a new recipient, Exchange Server sets the recipient's default email address based on the applicable email address policy.
- Whenever you apply an email address policy, Exchange Server automatically rewrites the email addresses for recipients to which the policy applies.

Journaling allows you to forward copies of messaging items and related reports automatically to an alternate location. You can use journaling to verify compliance with policies implemented in your organization and to help ensure that your organization can meet its legal and regulatory requirements. Enable journaling for the entire organization by using journal rules.

Transport rules, which apply to both on-premises and online Exchange organizations, allow you to screen messaging items and apply actions to those items that meet specific conditions. When you enable transport rules, all Mailbox servers in your Exchange organization screen messages according to the rules you've defined.

Creating and Managing Email Address Policies

Every Exchange organization has a default email address policy, which is required to create email addresses for recipients. You can create additional email address policies as well. For example, if your organization's internal domain name is different from its external domain name, you would need to create an accepted domain to match your external domain name and an email address policy that assigns your external domain name to user email addresses.

Working with Email Address Policies

To view the email address policies configured for your organization, complete the following steps:

1. In the Exchange Admin Center, select Mail Flow in the Navigation menu, and then select Email Address Policies.
2. In the main pane, email address policies are listed by name, priority, last modified time, and status as shown in Figure 23-1. The status is listed as Applied for a policy that has been applied to recipients and Unapplied for a policy that has not been applied to recipients.

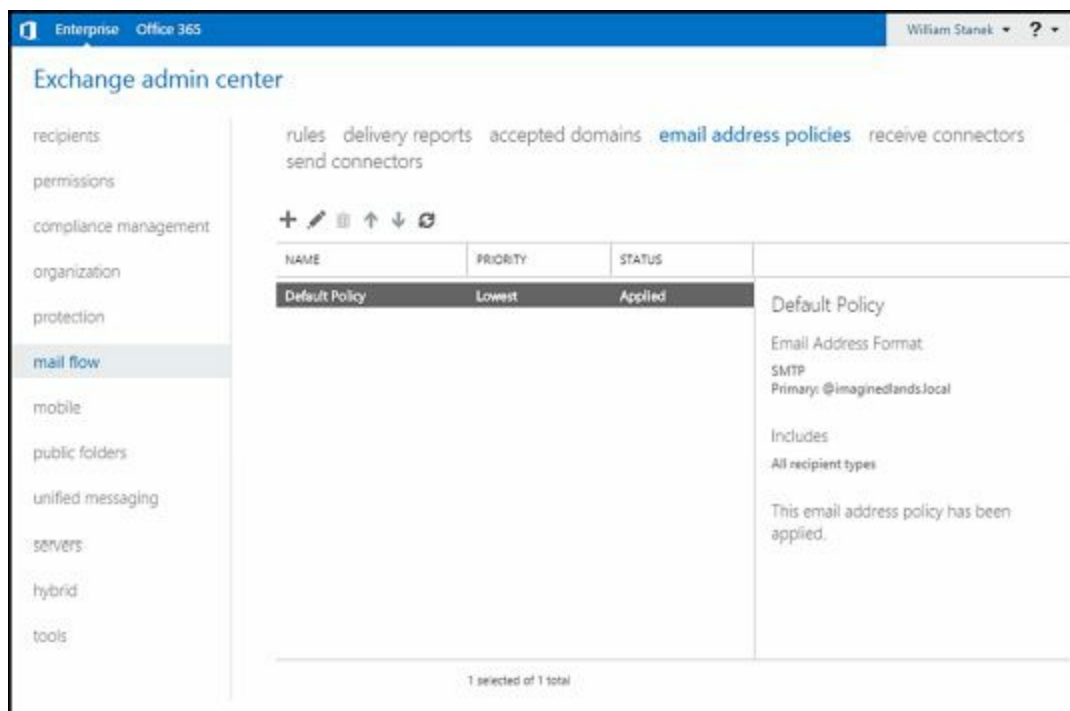


FIGURE 23-1 View the email address policies.

You can use the `Get-EmailAddressPolicy` cmdlet to list email address policies or to get information on a particular email address policy. If you don't provide an identity with this cmdlet, configuration information for all email address policies is displayed. Listing 23-1 provides the syntax and usage, as well as sample output, for the `Get-EmailAddressPolicy` cmdlet.

LISTING 23-1 Get-EmailAddressPolicy cmdlet syntax and usage

Syntax

```
Get-EmailAddressPolicy [-Identity PolicyIdentity]  
[-DomainController DCName] [-Organization OrgId]  
[-IncludeMailboxSettingOnlyPolicy <$true | $false>]
```

Usage

Get-EmailAddressPolicy | ft na me, priority, recipientfilter,
recipientfilterapplied, includedrecipients

```
Get-EmailAddressPolicy -Identity "Default Policy"
```

Output

Name	Priority	RecipientFilter
Default Policy	Lowest	Alias -ne \$null
Rewrite Group Addresses	1	Alias -ne \$null

Creating Email Address Policies

You can create email address policies for your organization by completing the following steps:

1. In the Exchange Admin Center, select Mail Flow in the Navigation menu, and then select Email Address Policies. Next, click New (**+**) to open the New Email Address Policy dialog box.

TYPE	ADDRESS FORMAT
SMTP	alias@imagedlands.com

2. Use the Name text box to identify the email address policy. You can use a descriptive name that identifies the purpose of the email address policy or simply enter the actual SMTP domain name to which it applies.
3. Under Email Address Format, click the Add button (**+**). This displays the Email Address Format dialog box.

4. Use the Email Address Format options to specify how to generate or rewrite email addresses automatically for each recipient to which the policy applies. You can use the Exchange alias or parts of the user name in various orders.
5. Use the Select An Accepted Domain drop-down list to select the email address domain. Only authoritative accepted domains are available for selection.
6. Although users can have multiple email addresses associated with their mailbox, only one email address, the default email address, is used for any sent messages. If you want the email address applied with this policy to be the default, select Make This Format The Reply Email Address.
7. Close the Email Address Format dialog box by clicking Save.
8. Specify the types of recipients to include in the policy. Select All Recipient Types, or select Only The Following Recipient Types, and then select the check boxes for the types of recipients to which you want to apply the policy.

9. If you've previously created other email address policies, set the relative priority of this policy. Policies are run in priority order. A policy with a priority of one has the highest priority and runs before a policy with a priority of two, and so on.
10. You can create rules that further filter recipients. Each rule acts as a condition that must be met. If you set more than one rule, each condition must be met for

there to be a match. To define a rule, click Add Rule. You can now set the filter conditions. The following types of conditions are available as well as conditions for custom attributes:

- **State Or Province** Filters recipients based on the value of the State/Province text box on the Contact Information page in the related Properties dialog box. In the Specify Words Or Phrases dialog box, type a state or province identifier to use as a filter condition, and then press Enter or click Add. Repeat as necessary, and then click OK.
- **Department** Filters recipients based on the value of the Department text box on the Organization page in the related Properties dialog box. In the Specify Words Or Phrases dialog box, type a department name to use as a filter condition, and then press Enter or click Add. Repeat as necessary, and then click OK.
- **Company** Filters recipients based on the value of the Company text box on the Organization page in the related Properties dialog box. In the Specify Words Or Phrases dialog box, type a company name to use as a filter condition, and then press Enter or click Add. Repeat as necessary, and then click OK.

IMPORTANT Although each rule acts as an OR condition for matches on specified values, the rules are aggregated as AND conditions. This means that a user that matches one of the values in a rule passes that filter but must be a match for all the rules to be included in the group. For example, if you were to define a state rule for Oregon, California, or Washington and a department rule for Technology, only users who are in Oregon, California, or Washington and in the Technology department match the filter.

11. Get a complete list of the recipients to which this policy will be applied by clicking Preview Recipients The Policy Applies To. If the policy applies to the expected recipients, click Save to create the email address policy. Otherwise, repeat Steps 4 to 11 and ensure you configure options and rules to appropriately define the recipients to which the policy should apply.
12. The policy is not applied automatically. To apply the policy, select the policy Exchange Admin Center's main pane, and then click Apply in the details pane.

If you click More Options in the Email Address Format dialog box, you'll be able to specify a custom SMTP email address. With custom addresses, you use the following variables to specify how the email address should be formatted as well as manually entered text:

- *%d* Inserts the recipient's display name
- *%g* Inserts the recipient's given name (first name)
- *%i* Inserts the recipient's middle initial
- *%m* Inserts the recipient's Exchange alias
- *%s* Inserts the recipient's surname (last name)
- *%ng* Inserts the first *N* letters of the given name.
- *%ns* Inserts the first *N* letters of the surname.

For example, you could enter %g.%s@tvpress.com to specify that email addresses should be formatted with the given name first, followed by a period (.) and the surname.

In Exchange Management Shell, you create and apply email address policies by using separate tasks. You can create email address policies by using the New-EmailAddressPolicy cmdlet. Once you create a policy, apply it using the Update-EmailAddressPolicy cmdlet. Listings 23-2 and 23-3 provide the syntax and usage for these cmdlets. Use the -EnabledPrimarySMTPAddressTemplate parameter to specify the custom format for email addresses. Although the syntax for custom email addresses is the same as when you are working with Exchange Admin Center, you must use the SMTP: prefix before specifying the format, as shown in the example.

NOTE Any time you receive an error regarding missing aliases, you should run the Update-EmailAddressPolicy cmdlet with the -FixMissingAlias parameter set to *True*. This tells Exchange to generate an alias for recipients who do not have an alias.

LISTING 23-2 New-EmailAddressPolicy cmdlet syntax and usage

Syntax

```
New-EmailAddressPolicy -Name PolicyName  
-EnabledPrimarySMTPAddressTemplate Template  
-IncludedRecipients RecipientTypes {AddtlParams} {ConditionalParams}
```

```
New-EmailAddressPolicy -Name PolicyName  
-EnabledEmailAddresses Templates -RecipientFilter Filter  
[-DisabledEmailAddresses Templates] {AddtlParams}
```

```
New-EmailAddressPolicy -Name PolicyName  
-EnabledPrimarySMTPAddressTemplate Template  
-RecipientFilter Filter {AddtlParams}
```

```
New-EmailAddressPolicy -Name PolicyName  
-EnabledEmailAddresses Templates  
-IncludedRecipients RecipientTypes  
[-DisabledEmailAddresses Templates]  
{AddtlParams} {ConditionalParams}
```

```
{AddtlParams}  
[-DomainController DCName] [-Organization OrgId ]  
[-Priority Priority ] [-RecipientContainer OUID ]
```

```
{ConditionalParams}  
[-ConditionalCompany CompanyNameFilter1 , CompanyNameFilter2,... ]  
[-ConditionalCustomAttribute N Value1, Value2, ...]  
[-ConditionalDepartment DeptNameFilter1 , DeptNameFilter2 , ... ]  
[-ConditionalStateOrProvince StateNameFilter1 , StateNameFilter2 , ... ]
```

Usage

```
New-EmailAddressPolicy -Name "Primary Email Address Policy"  
-IncludedRecipients "MailboxUsers, MailContacts, MailGroups"  
-ConditionalCompany "City Power & Light"  
-ConditionalDepartment "Sales", "Marketing"  
-ConditionalStateOrProvince "Washington", "Idaho", "Oregon"  
-Priority "Lowest"  
-EnabledEmailAddressTemplates "SMTP:%g.%s@tvpress.com"
```

LISTING 23-3 Update-EmailAddressPolicy cmdlet syntax and usage

Syntax

```
Update-EmailAddressPolicy -Identity PolicyIdentity  
[-DomainController DCName] [-FixMissingAlias <$true | $false>]
```

Usage


```
Update-EmailAddressPolicy -Identity "Primary Email Address Policy"
```

```
Update-EmailAddressPolicy -Identity "Primary Email Address Policy"  
-FixMissingAlias
```

Editing and Applying Email Address Policies

You can manage email address policies in several different ways. You can edit their properties or apply them to rewrite email addresses automatically for each recipient to which the policy applies. You can also change their priority to determine the precedence order for application in case there are conflicts between policies. When multiple policies apply to a recipient, the policy with the highest priority is the one that applies.

You can change the way email address policies work by completing the following steps:

1. In the Exchange Admin Center, select Mail Flow in the Navigation menu, and then select Email Address Policies.
2. In the main window, select the email address policy you want to change, and then select Edit (). This opens the properties dialog box for the policy.
3. Use the options on the General page to set the policy name and relative priority.
4. On the Email Address Format page, you can use the options provided to specify how to generate or rewrite email addresses automatically for each recipient to which the policy applies. You can use the Exchange alias or parts of the user name in various orders as described in steps 4 through 8 in the "Creating Email Address Policies" section of this chapter.
5. On the Apply To page, you can use the options provided to specify the recipients to which the policy will apply. After you configure options, preview the recipients to which the policy applies to ensure you've configured the settings appropriately, as described in steps 9 through 12 in the "Creating Email Address

Policies" section of this chapter.

6. The modified policy is not applied automatically. To apply the policy, select the policy Exchange Admin Center's main pane, and then click Apply in the details pane.

You can change priority in the Exchange Admin Center by selecting the policy, and then using the Increase Priority and Decrease Priority buttons to change the priority of the policy. The valid range for priorities depends on the number of policies you've configured. The Default Policy always has the lowest priority.

You can apply an email address policy by selecting the policy Exchange Admin Center's main pane, and then clicking Apply in the details pane.

In Exchange Management Shell, you can use the Set-EmailAddressPolicy cmdlet to modify email address policies, as shown in Listing 23-4. The Update-EmailAddressPolicy cmdlet, used to apply policies, was discussed previously.

LISTING 23-4 Set-EmailAddressPolicy cmdlet syntax and usage

Syntax

```
Set-EmailAddressPolicy -Identity PolicyIdentity
[-ConditionalCompany CompanyNameFilter1 , CompanyNameFilter2 ,... ]
[-ConditionalCustomAttribute N Value1 , Value2 , ...]
[-ConditionalDepartment DeptNameFilter1 , DeptNameFilter2 , ... ]
[-ConditionalStateOrProvince StateNameFilter1 , StateNameFilter2 , ... ]
[-DisabledEmailAddressTemplates Templates] [-DomainController DCName]
[-EnabledEmailAddressTemplates Templates]
[-EnabledPrimarySMTPAddressTemplate Template ]
[-ForceUpgrade <$true | $false>] [-IncludedRecipients RecipientTypes]
[-Name PolicyName] [-Priority Priority ]
[-RecipientContainer OUI] [-RecipientFilter Filter]
```


Usage

```
Set-EmailAddressPolicy -Identity "Primary Email Address Policy"
-Name "Tvpress.com Email Address Policy"
-IncludedRecipients "MailboxUsers"
-ConditionalCompany "City Power & Light"
-ConditionalDepartment "Sales"
-ConditionalStateOrProvince "Washington"
-Priority "2"
-EnabledEmailAddressTemplates "SMTP:%g.%s@tvpress.com"
```

Removing Email Address Policies

You can remove an email address policy that is no longer needed by completing the following steps:

1. In the Exchange Admin Center, select Mail Flow in the Navigation menu, and then select Email Address Policies.

2. In the main window, select the email address policy you want to delete, and then select Delete ().
3. When prompted to confirm, click Yes.

In Exchange Management Shell, you can use the `Remove-EmailAddressPolicy` cmdlet to remove email address policies. Listing 23-5 provides the syntax and usage.

LISTING 23-5 Remove-EmailAddressPolicy cmdlet syntax and usage

Syntax

```
Remove-EmailAddressPolicy -Identity EmailAddressPolicyIdentity  
[-DomainController DCName]
```

Usage

```
Remove-EmailAddressPolicy -Identity "Tvpress.com  
Email Address Policy"
```

Configuring Journal Rules

Exchange 2016 Setup creates a separate container in Active Directory Domain Services to store Exchange 2016 journal rules. If you are installing Exchange 2016 in an existing Exchange 2010 or Exchange 2013 organization, Setup copies any existing journal rules to this container so they will be applied to Exchange 2016. If you subsequently make changes to the journal rule configuration on your Exchange 2010 or Exchange 2013 servers, you must make the same changes on Exchange 2016 to ensure the journal rules are consistent across the organization (and vice versa). You can also export journal rules from Exchange 2010 or Exchange 2013 and import them to Exchange 2016.

Both Exchange Online and on-premises Exchange support a full set of compliance options for in-place eDiscovery and hold, auditing, retention policies, retention tags, and journal rules. These compliance options are configured in much the same way whether you are working with Exchange Online or Exchange 2016. If you are working in a hybrid configuration and have specific compliance requirements, you can ensure your on-premises compliance settings are applied to Exchange Online.

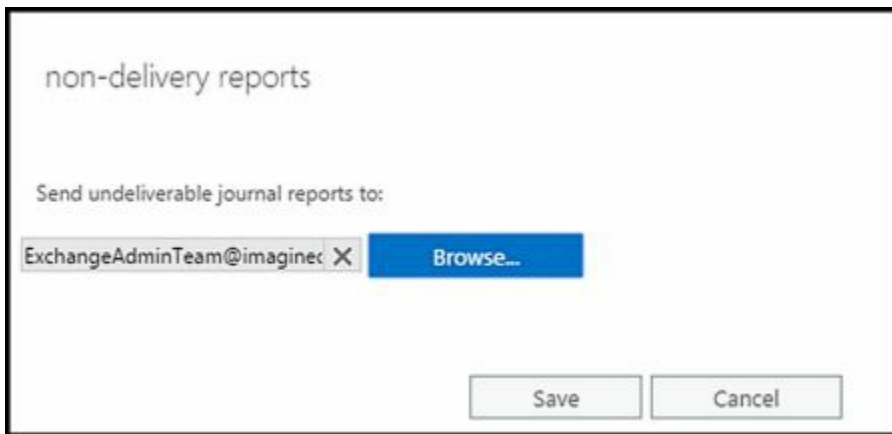
In a hybrid environment, inbound and outbound messages have separate routing configurations. If you enable centralized mail transport for inbound, outbound, or both types of messages in the hybrid configuration, messages sent by or to recipients in the online organization are set through the on-premises organization to ensure that compliance rules and any other processes or messaging requirements configured in the on-premises organization are applied. However, there is a noteworthy exception: Outbound messages sent from Exchange Online to other recipients in the same Exchange Online organization are delivered directly.

Setting The NDR Journaling Mailbox

When you first configure journaling, you'll need to specify an email address to receive any journal reports that are otherwise undeliverable. Typically, you'll want to create a new, dedicated mailbox to receive these reports so that the mailbox will not be journaled and also won't be subject to any transport rules or mailbox rule settings.

To specify the NDR journaling mailbox, complete the following steps:

1. In the Exchange Admin Center, select Compliance Management in the Navigation menu, and then select Journal Rules.
2. If the journaling mailbox has already been specified, the email address is listed; otherwise, click Select Address.



3. In the NDRs For Undeliverable Journal Reports dialog box, click Browse, select a destination mailbox, and then click OK.
4. Click Save.
5. A warning prompts lets you know the email address specified won't be journaled or use transport/mailbox rules. Click OK.



Creating Journal Rules

You create journal rules to record messages in your organization in support of email retention and compliance requirements. You can target journal rules for the following:

- **Internal messaging items** Tracks messaging items sent and received by recipients inside your Exchange organization.
- **External messaging items** Tracks messaging items sent to recipients or from senders outside your Exchange organization.
- **Allmessaging items** Tracks all messaging items, including those already processed by journal rules that track only internal or external messaging items.

When you enable journal rules for one or more of these scopes, the rules are executed on your organization's Mailbox servers. Journal rules can be targeted to all recipients or to specific recipients. For example, you can create a rule to journal all messages sent to the AllEmployees distribution group.

You can create a journal rule by completing the following steps:

1. In the Exchange Admin Center, select Compliance Management in the Navigation

menu, and then select Journal Rules. Next, click New to open the New Journal Rule dialog box (shown in Figure 23-2).

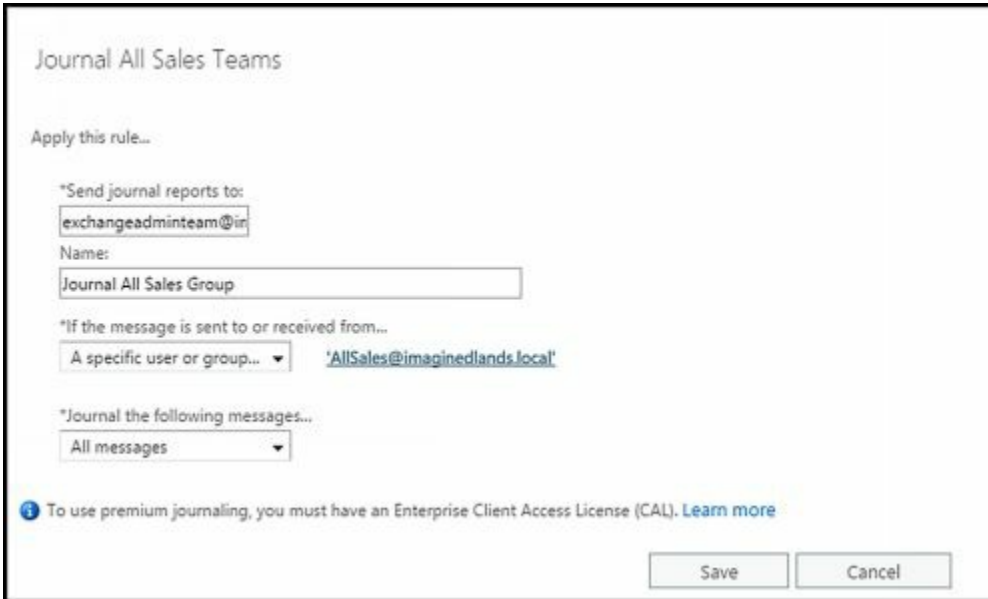
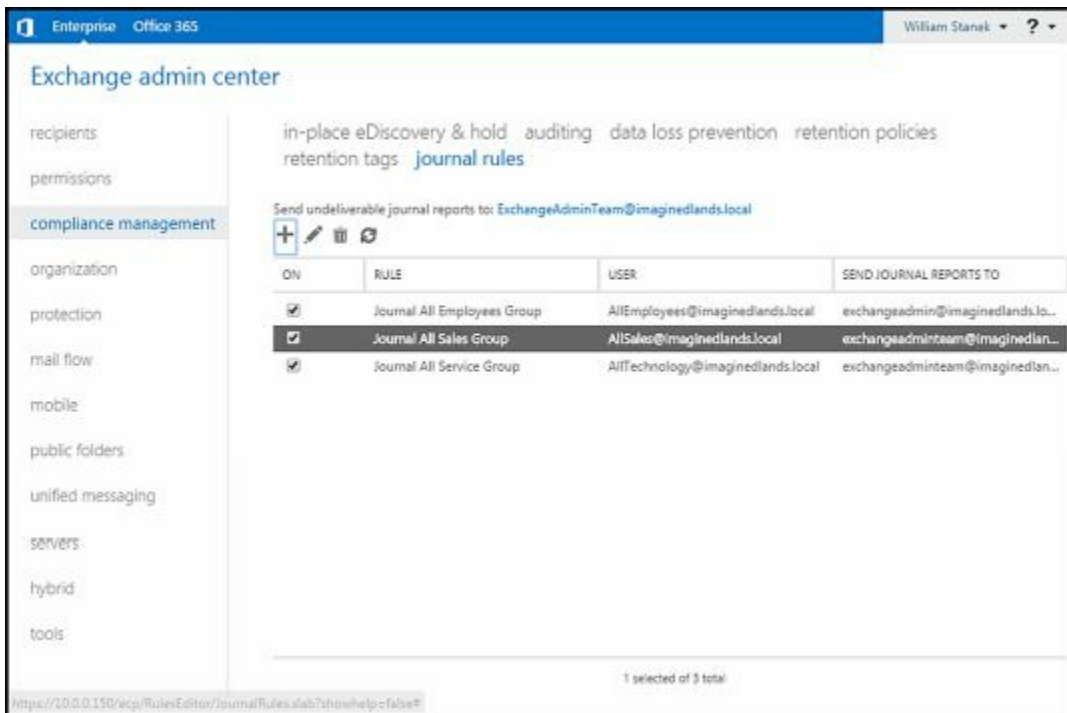




FIGURE 23-2 Create a journal rule.

2. In the Name text box, type a descriptive name for the rule.
3. In the Send Journal Reports To, provide the journal email address. This is the recipient to which Exchange should forward journal reports for this rule.
4. Use the If The Message Is Send To Or Received From selection list to specify whether the rule should be applied to messages sent to or received from a specific user or group, or to all messages. For a specific user or group, you'll then need to select the user or group.
5. Use the Journal The Following Messages selection list to specify the scope of the rule as either All Messages, Internal Messages Only, or External Messages Only.
6. Click Save.

Managing Journal Rules

When you are working with the Compliance Management area and select Journal Rules, the currently defined journal rules are listed in the main pane by status, name, user journaled, and journal report recipient.



You can easily enable or disable a rule by setting or clearing the corresponding checkbox in the On column. If you select a rule and then select Edit (), you can modify the rule settings. If you select a rule and then select Delete (), you can delete the rule.

In Exchange Management Shell, you can manage journal rules by using the following cmdlets: New-JournalRule, Set-JournalRule, Get-JournalRule, and Remove-JournalRule.

Configuring Transport Rules

Transport rules have conditions, actions, and exceptions that you can apply. Conditions you can screen for include the following:

- **The sender is...** Allows you to screen messages from specific senders according to their email address, group membership, account properties and more.
- **The recipient is...** Allows you to screen messages being sent to specific recipients according to their email address, group membership, account properties, and so forth.
- **The subject or body...** Allows you to screen messages that have specific words in their subject line or message body.
- **Any attachment...** Allows you to screen messages with attachments for specific content, file names, file extensions, and password protection as well as when the attachment size is greater than or equal to the size limit that you set.
- **The message...** Allows you to screen messages sent to or copied to specific recipients or specific groups as well as when the message size is greater than or equal to a size limit that you set.
- **The send and recipient...** Allows you to screen messages sent between members of specific groups, messages sent by subordinates of a specific manager, and messages sent to a recipient who is a manager or direct report of the sender.
- **The message properties...** Allows you to screen messages that have a spam confidence level (SCL) rating that is greater than or equal to a limit that you set. Also allows you to screen messages by classification, type, and importance level.
- **A message header...** Allows you to screen messages that have a header field that includes specific words or matches specific patterns.

When a message meets all of the conditions you specify in a transport rule, the message is handled according to the actions you've defined. Actions you can apply to messages that meet your transport rule conditions include the following:

- Forwarding the message for approval to specific people or to the sender's manager.
- Redirecting the message to specific recipients, host quarantine or a specific outbound connector.
- Blocking the message by rejecting it with a specific return message and explanation or by deleting the message without notifying anyone.
- Adding recipients to the Bcc, To, or Cc fields or simply adding the sender's manager as a recipient.
- Applying a disclaimer to the beginning or end of the message.
- Modifying the message by removing a message header, adding a message header, adding a message classification, or setting the spam confidence level.
- Securing the message with rights protection or TLS encryption.
- Prepending the subject of the message with a specified text value.
- Generating an incident report and sending it to specific recipients.

Transport rules can also have exceptions. Exception criteria are similar to condition

criteria. For example, you can exclude messages from certain people or from certain members of distribution lists. You can also exclude messages sent to certain people or to particular members of a distribution list.

Creating Transport Rules

You can create a transport rule by completing the following steps:


1. In the Exchange Admin Center, select Mail Flow in the Feature area, and then select Rules. Next, click New (), and then select Create A New Rule to open the New Rule dialog box.
2. In the New Rule dialog box, click the More Option link to display additional options.

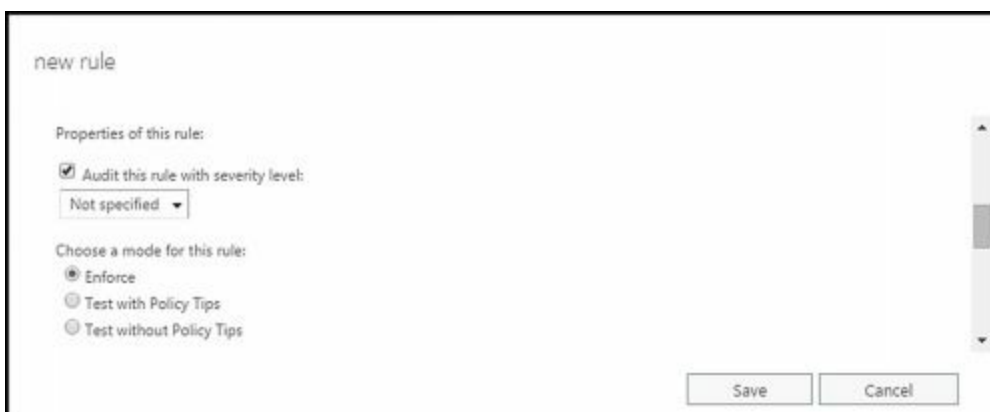


FIGURE 23-3 Create a transport rule.

3. In the Name text box, type a descriptive name for the rule and optionally enter a descriptive comment.
4. Next you need to specify the conditions for the rule by using the options under Apply This Rule If... Click in the selection list. Next, choose what part of the message to examine and then choose how the condition should be matched. Finally, in the selection dialog box, set the condition parameters by selecting an option or typing a value or values to match. For example, choose "The sender" and then choose "is this person." Finally, in the Select Members dialog box, select the sender to match in the rule and then click OK.
5. If you want to add another condition, click Add Condition and then click in the new selection list provided. Next, choose what part of the message to examine, and then choose how the condition should be matched. Finally, in the selection dialog box, set the condition parameters by selecting an option or typing a value or values to match. Repeat this step to add other conditions.
6. Use the options under Do The Following... to define the actions to take when a message meets the condition or conditions you specified. Click in the selection

list. Next, choose the action, and then choose how the action should be performed. Finally, in the selection dialog box, set the action parameters by selecting an option or typing a value or values to match. For example, choose "Add recipients" and then choose "to the Bcc box." Finally, in the Select Members dialog box, select the sender that should be added to the Bcc field of matching messages and then click OK.

7. If you want to add another action, click Add Action, and then click in the new selection list provided. Next, choose the action, and then choose how the action should be performed. Finally, in the selection dialog box, set the action parameters by selecting an option or typing a value or values to match. Repeat this step to add other actions.
8. You now need to specify any exceptions by using the options under Except If... Click Add Exception, and then click in the new selection list provided. Next, choose the exception and then choose how the exception should be matched. Finally, in the selection dialog box, set the exception parameters by selecting an option or typing a value or values to match. For example, choose "The sender" and then choose "is this person." Finally, in the Select Members dialog box, select the sender to add as an exception to the rule, and then click OK. Repeat this step to add other exceptions.
9. By default, transport rules are audited and enforced. If you don't want the rule to be audited, clear the Audit This Rule... check box. If you want to test the rule rather than enforce it, select Test With Policy Tips or Test Without Policy Tips instead of Enforce.





10. Click Save to create the rule. If an error occurs during rule creation, note the error and then correct the issue before trying to create the rule again.

Managing Transport Rules

You can manage transport rules in several different ways including editing their properties or disabling them. When you've created multiple rules, you can also change their priority to determine the precedence order for application in case there are conflicts between rules. When multiple rules apply to a message, the rule with the highest priority is the one that your Mailbox server applies.

When you select a transport rule in Exchange Admin Center, you can manage the rule in

the following ways:

- **Change Priority** Use the Move Up or Move Down buttons to increase or decrease the relative priority of the rule. Rules are processed in priority order, with the rule listed first being the first one processed, the rule listed second being the second processed, and so on.
- **Disable Rule** Use the checkbox in the On column to enable or disable the rule.
- **Remove** Select Delete () to remove the rule.
- **Edit Rule** Select Edit () to edit the properties of the transport rule.

In Exchange Management Shell, you can manage transport rules by using the following cmdlets: New-TransportRule, Set-TransportRule, Get-TransportRule, and Remove-TransportRule.

Chapter 24. Filtering Spam

Every minute users spend dealing with unsolicited commercial email (spam) or other unwanted email is a minute they cannot do their work and address other issues. To try to deter spammers and other unwanted senders, you can use message filtering to block these senders from directing messages to your organization. Not only can you filter messages that claim to be from a particular sender or that are sent to a particular receiver, you can also establish connection filtering rules based on IP block lists. The sections that follow discuss these and other anti-spam options.

As you configure message filtering, keep in mind that although Exchange Server is designed to combat most spammer techniques, no system can block all of them. Like the techniques of those who create viruses, the techniques of those who send spam frequently change, and you won't be able to prevent all unwanted email from going through. You should, however, be able to substantially reduce the flow of spam into your organization.

The way you configure message filtering depends on whether your organization uses Edge Transport servers:

- If your organization doesn't use Edge Transport servers and has only Mailbox servers, you can enable the anti-spam features on Mailbox servers that receive messages from the Internet and then configure message filtering options on those servers.
- If your organization is using Edge Transports, you can enable the anti-spam features on the Edge Transport servers and then and configure message filtering on those servers.

As discussed in "Enabling Anti-Spam Features" in Chapter 21 "Configuring Transport Services," Edge Transport servers have anti-spam features enabled by default and Mailbox servers do not. Generally, you want your Edge Transports to handle spam filtering before messages are routed into the Exchange organization. After your Edge Transport servers have filtered messages, there is no need to filter them again, which is why message filtering is disabled by default on Mailbox servers.

Filtering Spam by Sender

Sometimes, when you are filtering spam or other unwanted email, you'll know specific email addresses or email domains from which you don't want to accept messages. In this case, you can block messages from these senders or email domains by configuring sender filtering. Another sender from which you probably don't want to accept messages is a blank sender. If the sender is blank, it means the From field of the email message wasn't filled in and the message is probably from a spammer.

Sender filtering is enabled by default and is designed to filter inbound messages from non-authenticated Internet sources. You can view the current configuration of sender filtering by using Get-SenderFilterConfig. Use the -Enabled parameter of Set-SenderFilterConfig to enable or disable sender filtering. The following example disables sender filtering:

```
Set-SenderFilterConfig -Enabled $false
```

In Exchange Management Shell, you can use the Set-SenderFilterConfig cmdlet to configure sender filtering.

Listing 24-1 provides the syntax and usage. You can block individual email addresses using the -BlockedSenders parameters. If you want to filter all email sent from a particular domain, use the -BlockedDomains or -BlockedDomainsAndSubdomains parameter.

LISTING 24-1 Set-SenderFilterConfig cmdlet syntax and usage

Syntax

```
Set-SenderFilterConfig [-Action <StampStatus | Reject>]  
[-BlankSenderBlockingEnabled <$true | $false>]  
[-BlockedDomains <domain1, domain2...domainN>]  
[-BlockedDomainsAndSubdomains <domain1, domain2...domainN>]  
[-BlockedSenders <email1, email2...emailN>]  
[-DomainController DCName]  
[-Enabled <$true | $false>]  
[-ExternalMailEnabled <$true | $false>]  
[-InternalMailEnabled <$true | $false>]  
[-RecipientBlockedSenderAction <Reject | Delete>]
```

Usage

```
Set-SenderFilterConfig -BlankSenderBlockingEnabled $true
```

```
Set-SenderFilterConfig -BlockedDomains contoso.com, margiestravel.com,  
proseware.com
```

```
Set-SenderFilterConfig -BlockedDomainsAndSubdomains fineartschool.com,  
wingtiptoys.com
```

```
Set-SenderFilterConfig -BlockedSenders tony@treymresearch.net,  
ed@woodgrovebank.com
```

By default, sender filtering rejects messages from blocked domains and senders. This option ensures that Exchange doesn't waste processing power and other resources dealing with messages from filtered senders. If you want to mark messages as being from a blocked sender and continue processing them, set the `-Action` parameter to `StampStatus` instead. Here, a message header stamp will be added to the message and the message will be processed by other anti-spam agents. This stamp and any other issues found will then be used to set the spam confidence level as part of content filtering. A message that exceeds a spam confidence level is rejected, quarantined, deleted, or marked as junk mail. Set the `-BlankSenderBlockingEnabled` parameter to `$true` to block blank senders.

As shown in the previous examples, you can easily define the initial set of blocked domains and senders. If you want to modify these values, however, you must either enter the complete set of blocked domains or senders, or you must use a special shorthand to insert into or remove values from these multivalued properties. The shorthand for adding values is:

```
@{Add="<ValuetoAdd1>","<ValuetoAdd2>"...}
```

Such as:

```
Set-SenderFilterConfig -BlockedDomains @{Add="imaginedlands.com","tvpress.com"}
```

The shorthand for removing values is:

```
@{Remove="<ValuetoRemove1>","<ValuetoRemove2>"...}
```

Such as:

```
Set-SenderFilterConfig -BlockedDomains  
@{Remove="imaginedlands.com","tvpress.com"}
```

If you want to add values and remove others, you can do this as well by using the following shorthand:

```
@{Add="<ValuetoAdd1>","<ValuetoAdd2>"...;  
Remove="<ValuetoRemove1>","<ValuetoRemove1>"...}
```

You can confirm that values were added or removed as expected by using `Get-SenderFilterConfig`. In this example, you view the currently blocked domains:

```
Get-SenderFilterConfig | fl BlockedDomains
```

By default, `-InternalMailEnabled` is set to `$false` and `-ExternalMailEnabled` is set to `true` which means authenticated internal email messages aren't processed by the Sender Filter whereas unauthenticated external email messages are processed by the Sender filter. An unauthenticated external email messages is one from an untrusted or anonymous source rather than a trusted partner.

Finally, when users in your organization add senders to their blocked sender list, the SafeList aggregation feature in Exchange 2016 adds these senders to the Blocked Senders List in Exchange 2016. By default, messages from these users are rejected rather than deleted. To delete these messages, set -RecipientBlockedSenderAction to Delete.

Filtering Spam by Recipient

In any organization, you'll have users whose email addresses change, perhaps because they request it, leave the company, or change office locations. Although you might be able to forward email to these users for a time, you probably won't want to forward email indefinitely. At some point, you, or someone else in the organization, will decide it's time to delete the user's account, mailbox, or both. If the user is subscribed to mailing lists or other services that deliver automated email, the automated messages continue to come in, unless you manually unsubscribe the user or reply to each email that you don't want to receive the messages. Unfortunately, some Exchange administrators find themselves going through this inefficient process. It's much easier to add the old or invalid email address to a recipient filter list and specify that Exchange shouldn't accept messages for users who aren't in the Exchange directory. Once you do this, Exchange won't attempt to deliver messages for filtered or invalid recipients, and you won't see related nondelivery reports (NDRs).

Recipient filtering is enabled by default. In Exchange Management Shell, you can use the `Set-RecipientFilterConfig` cmdlet to configure recipient filtering. Listing 24-2 provides the syntax and usage.

LISTING 24-2 Set-RecipientFilterConfig cmdlet syntax and usage

Syntax

```
Set-RecipientFilterConfig [-BlockedRecipients <email1,email2...emailN>]
[-BlockListEnabled <$true | $false>] [-DomainController DCName]
[-Enabled <$true | $false>] [-ExternalMailEnabled <$true | $false>]
[-InternalMailEnabled <$true | $false>]
[-RecipientValidationEnabled <$true | $false>]
```

Usage

```
Set-RecipientFilterConfig -RecipientValidationEnabled $true
```

```
Set-RecipientFilterConfig -BlockedRecipients tony@treyresearch.net,
ed@woodgrovebank.com
```

By default, recipient filtering rejects messages from blocked recipients but doesn't block users from sending messages to blocked recipients. If you set `-BlockListEnabled` to `$true`, users won't be able to send messages to blocked recipients. You also can specify whether Exchange 2016 validates recipients and then blocks messages sent to recipients who don't exist. Although Exchange 2016 doesn't validate recipients by default, you can have Exchange 2016 validate recipients by setting `-RecipientValidationEnabled` to `$true`.

Blocked list addresses can refer to a specific email address, such as `walter@blueyonderairlines.com`, or a group of email addresses designated with the wildcard character (`*`), such as `*@blueyonderairlines.com` to filter all email addresses

from blueyonderairlines.com, or *@*.blueyonderairlines.com, to filter all email addresses from child domains of blueyonderairlines.com.

If you want to modify the blocked recipients, you must either enter the complete set of blocked recipients, or you use a special shorthand to insert into or remove values from this multivalued property. The shorthand for adding values is:

```
@{Add="<ValuetoAdd1>","<ValuetoAdd2>"...}
```

Such as:

```
Set-RecipientFilterConfig -BlockedRecipients
```

```
@{Add="mary@imaginedlands.com","gene@tvpress.com"}
```

The shorthand for removing values is:

```
@{Remove="<ValuetoRemove1>","<ValuetoRemove2>"...}
```

Such as:

```
Set-RecipientFilterConfig -BlockedRecipients
```

```
@{Remove="mary@imaginedlands.com","gene@tvpress.com"}
```

If you want to add values and remove others, you can do this as well by using the following shorthand:

```
@{Add="<ValuetoAdd1>","<ValuetoAdd2>"...;  
Remove="<ValuetoRemove1>","<ValuetoRemove1>"...}
```

You can confirm that values were added or removed as expected by using Get-RecipientFilterConfig. In this example, you view the currently blocked domains:

```
Get-RecipientFilterConfig | fl BlockedRecipients
```

Filtering Connections with IP Block Lists

If you find that sender and recipient filtering isn't enough to stem the flow of spam into your organization, you might want to consider subscribing to an IP block list service. Here's how this service works:

- You subscribe to an IP block list service. Although there are free services available, you might have to pay a monthly service fee. In return, the service lets you query its servers for known sources of unsolicited email and known relay servers.
- The service provides you with domains you can use for validation and a list of status codes to watch for. You configure Exchange to use the specified domains and enter connection filtering rules to match the return codes, and then you configure any exceptions for recipient email addresses or sender IP addresses.
- Each time an incoming connection is made, Exchange performs a lookup of the source IP address in the block list domain. A “host not found” error is returned to indicate the IP address is not on the block list and no match was found. If there is a match, the block list service returns a status code that indicates the suspected activity. For example, a status code of 127.0.0.3 might mean that the IP address is from a known source of unsolicited email.
- If there is a match between the status code returned and the filtering rules you've configured, Exchange returns an error message to the user or server attempting to make the connection. The default error message says that the IP address has been blocked by a connection filter rule, but you can specify a custom error message to return instead.

The sections that follow discuss applying IP block lists, setting provider priority, defining custom error messages to return, and configuring block list exceptions. These tasks will need to be performed when you work with IP block lists.

Applying IP Block Lists

Before you get started, you need to know the domain of the block list service provider, and you should also consider how you want to handle the status codes the provider returns. Exchange allows you to specify that any return status code is a match, that only a specific code matched to a bit mask is a match, or that any of several status codes that you designate can match.

A list of typical status codes that might be returned by a provider service, include:

- 127.0.0.1 Trusted nonspam (on the “white” list). Code bit mask: 0.0.0.1.
- 127.0.0.2 Known source of unsolicited email/spam (on the “black” list). Code bit mask: 0.0.0.2.
- 127.0.0.3 Possible spam, like a mix of spam and nonspam (on the “yellow” list). Code bit mask: 0.0.0.3.
- 127.0.0.4 Known source of unsolicited email/spam, but not yet blocked (on the “brown” list). Code bit mask: 0.0.0.4.

- 127.0.0.5 Not a spam-only source, and not on the “black” list. Code bit mask: 0.0.0.5.

Rather than filter all return codes, in most cases, you’ll want to be as specific as possible about the types of status codes that match to help ensure that you don’t accidentally filter valid email. For example, based on the list of status codes of the provider, you might decide that you want to filter known sources of unsolicited email and known relay servers, but not filter known sources of dial-up user accounts, which might or might not be sources of unsolicited email.

Exchange 2016 allows you to configure multiple IP block list providers, with a relative priority assigned to a provider determining the order in which providers are checked. In Exchange Management Shell, you manage IP block list providers and their settings by using the following:

- **Add-IPBlockListProvider** Adds an IP block list provider.

```
Add-IPBlockListProvider -LookupDomain SmtpDomain -Name ProviderName
[-AnyMatch <$true | $false>] [-BitmaskMatch IPAddressBitMask]
[-DomainController DCName] [-Enabled <$true | $false>]
[-IPAddressesMatch IpAddress1,IpAddress2...IpAddressN]
[-Priority Priority] [-RejectionResponse Response]
```

- **Get-IPBlockListProvider** Displays the settings of a specific or all IP block list providers.

```
Get-IPBlockListProvider [-Identity SmtpDomain]
[-DomainController DCName]
```

- **Set-IPBlockListProvider** Modifies the settings associated with the specified IP block list provider.

```
Set-IPBlockListProvider -Identity SmtpDomain
[-AnyMatch <$true | $false>] [-BitmaskMatch IPAddressBitMask]
[-DomainController DCName] [-Enabled <$true | $false>]
[-IPAddressesMatch IpAddress1,IpAddress2...IpAddressN]
[-Priority Priority] [-RejectionResponse Response]
```

- **Remove-IPBlockListProvider** Removes a IP block list provider.

```
Remove-IPBlockListProvider -Identity SmtpDomain
[-DomainController DCName]
```

When you add a block list provider, you use the -Name parameter to set a descriptive name for the provider and the -LookupDomain to specify the domain name of the block list provider service, such as *proseware.com*. You can then specify whether to match any return code (other than an error) received from the provider service or to match a specific mask and return codes from the provider service. As shown in the following example, set -AnyMatch to \$true to match any return code:

```
Add-IPBlockListProvider -Name Proseware -LookupDomain proseware.com
-AnyMatch $true
```

If you want to match a specific mask, use `-BitmaskMatch` to specify the bitmask to match, such as:

```
Add-IPBlockListProvider -Name Proseware -LookupDomain proseware.com  
-BitmaskMatch 0.0.0.4
```

Alternatively, you can match specific values in the return status codes by using `-IPAddressesMatch`, such as:

```
Add-IPBlockListProvider -Name Proseware -LookupDomain proseware.com  
-IPAddressesMatch 127.0.0.4, 127.0.0.5, 127.0.0.6, 127.0.0.7
```

Other commands you can use to manage and work with block lists include:

- [Get-IPBlockListConfig](#) Displays information about the configuration of the Connection Filter agent.

```
Get-IPBlockListConfig [-DomainController DCName]
```

- [Set-IPBlockListConfig](#) Modifies the configuration of the Connection Filter agent.

```
Set-IPBlockListConfig [-DomainController DCName]  
[-Enabled <$true | $false>] [-ExternalMailEnabled <$true | $false>]  
[-InternalMailEnabled <$true | $false>]  
[-MachineEntryRejectionResponse Response]  
[-StaticEntryRejectionResponse Response]
```

- [Set-IPBlockListProvidersConfig](#) Modifies the block list provider configuration used by Connection Filter agent.

```
Set-IPBlockListProvidersConfig [-DomainController DCName]  
[-BypassedRecipients <email1,email2...emailN>]  
[-Enabled <$true | $false>]  
[-ExternalMailEnabled <$true | $false>]  
[-InternalMailEnabled <$true | $false>]
```

- [Test-IPBlockListProvider](#) Checks connectivity to the specified block list provider and then issues a lookup request for an IP address to verify.

```
Test-IPBlockListProvider -Identity SmtPName -IPAddress IPAddress  
[-DomainController DCName] [-Server ServerID ]
```

Configuring Block List Providers

You can configure multiple block list providers. Each provider is listed in priority order, and if Exchange makes a match by using a particular provider, the other providers are not checked for possible matches. In addition to being prioritized, providers can also be enabled or disabled. If you disable a provider, it's ignored when looking for possible status code matches.

In Exchange Management Shell, you can use `Add-IPBlockListProvider` and `Set-IPBlockListProvider` to manage provider priority and to enable or disable providers. If you don't specify a priority when you add a provider using `Add-IPBlockListProvider`, the order providers are added sets the priority, with the first provider added having a

priority of 1, the second a priority of 2, and so on.

Use the `-Priority` parameter to set the relative priority of a provider and the `-Enabled` parameter to enable or disable a provider. In this example, you set the priority of `Proseware.com` to 2:

```
Add-IPBlockListProvider -Identity Proseware.com -Priority 2
```

Specifying Custom Error Messages

When a match is made between the status code returned and the filtering rules you've configured for block list providers, Exchange returns an error message to the user or server attempting to make the connection. The default error message says that the IP address has been blocked by a connection filter rule. If you want to override the default error message, you can specify a custom error message to return on a per-rule basis. The error message can contain the following substitution values:

- `%0` to insert the connecting IP address
- `%1` to insert the name of the connection filter rule
- `%2` to insert the domain name of the block list provider service

Some examples of custom error messages include the following:

- The IP address (`%0`) was blocked and not allowed to connect.
- `%0` was rejected by `%2` as a potential source of unsolicited email.

The custom error message can't be more than 240 characters.

In Exchange Management Shell, you use the `-RejectionResponse` parameter of `Add-IPBlockListProvider` and `Set-IPBlockListProvider` to set a custom error message on a per-provider basis. Use the `-RejectionResponse` parameter with the `Set-ContentFilterConfig` cmdlet to set the default custom response.

Defining Block Lists

Sometimes, you'll find that an IP address, a network, or an email address shows up incorrectly on a block list. The easiest way to correct this problem is to create a block list exception that indicates that the specific IP address, network, or email address shouldn't be filtered.

Using Connection Filter Exceptions

You can define connection filter exceptions for email addresses using the `-BypassedRecipients` parameter of the `Set-IPBlockListProvidersConfig` cmdlet. Define the initial set of exceptions simply by entering the email addresses in a comma-separated list, such as:

```
Set-IPBlockListProvidersConfig -BypassedRecipients joe@imaginedlands.com,  
sarah@tvpress.com
```

If you want to modify the exceptions, however, you must either enter the complete set of exceptions, or use a special shorthand to insert into or remove values from this multivalued property. The shorthand for working with multivalued properties is:

```
@{Add="<ValuetoAdd1>","<ValuetoAdd2>"...}
```

```
@{Remove="<ValuetoRemove1>","<ValuetoRemove2>"...}
```

```
@{Add="<ValuetoAdd1>","<ValuetoAdd2>"...;  
Remove="<ValuetoRemove1>","<ValuetoRemove1>"...}
```

Such as:

```
Set-IPBlockListProvidersConfig -BypassedRecipients  
@{Add="tina@tresearch.net","mark@contosso.com";  
Remove="sarah@tvpress.com"}
```

Using Global Allowed Lists

Exchange will accept email from any IP address or network on the global allowed list. Before you can define allowed entries for IP addresses and networks you must be sure that the IP Allow List is enabled.

You use `Add-IPAllowListEntry` to add an IP address or IP address range to the IP Allow list. Listing 24-3 provides the syntax and usage.

LISTING 24-3 Add-IPAllowListEntry cmdlet syntax and usage

Syntax

```
Add-IPAllowListEntry -IPAddress IPAddress {AddtlParams}
```

```
Add-IPAllowListEntry -IPRange IPRange {AddtlParams}
```

```
{AddtlParams}  
[-Comment Comment] [-ExpirationTime DateTime] [-Server ServerId]
```

Usage

```
Add-IPAllowListEntry -IPAddress 192.168.10.45
```

```
Add-IPAllowListEntry -IPRange 192.168.10.0/24
```

```
Add-IPAllowListEntry -IPRange 192.168.10.1-192.168.10.254
```

You use `Get-IPAllowListEntry` to list IP Allow List entries and `Remove-IPAllowListEntry` to remove IP Allow List entries. Listings 24-4 and 24-5 provide the syntax and usage.

LISTING 24-4 `Get-IPAllowListEntry` cmdlet syntax and usage

Syntax

```
Get-IPAllowListEntry [-Identity IPListEntryId] {AddtlParams}
```

```
Get-IPAllowListEntry -IPAddress IPAddress {AddtlParams}
```

```
{AddtlParams}  
[-ResultSize Size] [-Server ServerId]
```

Usage

```
Get-IPAllowListEntry
```

```
Get-IPAllowListEntry -IPAddress 192.168.10.45
```

LISTING 24-5 `Remove-IPAllowListEntry` cmdlet syntax and usage

Syntax

```
Remove-IPAllowListEntry -Identity IPListEntryId  
[-Server ServerId]
```

Usage

```
Get-IPAllowListEntry | Where {$_.IPRange -eq '192.168.10.45'} |  
Remove-IPAllowListEntry
```

```
Get-IPAllowListEntry | Where {$_.IPRange -eq '192.168.10.0/24'} |  
Remove-IPAllowListEntry
```

Using Global Block Lists

Exchange will reject email from any IP address or network on the block list. Before you can define blocked entries for IP addresses and networks, you must ensure that the IP block list is enabled.

You use `Add-IPBlockListEntry` to add an IP address or IP address range to the IP block list. Listing 24-6 provides the syntax and usage.

LISTING 24-6 Add-IPBlockListEntry cmdlet syntax and usage

Syntax

```
Add-IPBlockListEntry -IPAddress IPAddress {AddtlParams}
```

```
Add-IPBlockListEntry -IPRange IPRange {AddtlParams}
```

```
{AddtlParams}
```

```
[-Comment Comment] [-ExpirationTime DateTime ] [-Server ServerId]
```

Usage

```
Add-IPBlockListEntry -IPAddress 192.168.10.45
```

```
Add-IPBlockListEntry -IPRange 192.168.10.0/24
```

```
Add-IPBlockListEntry -IPRange 192.168.10.1-192.168.10.254
```

You use Get-IPBlockListEntry to list IP block list entries and Remove-IPBlockListEntry to remove IP block list entries. Listings 24-7 and 24-8 provide the syntax and usage.

LISTING 24-7 Get-IPBlockListEntry cmdlet syntax and usage

Syntax

```
Get-IPBlockListEntry [-Identity IPListEntryId] {AddtlParams}
```

```
Get-IPBlockListEntry -IPAddress IPAddress {AddtlParams}
```

```
{AddtlParams}
```

```
[-ResultSize Size ] [-Server ServerId]
```

Usage

```
Get-IPBlockListEntry
```

```
Get-IPBlockListEntry -IPAddress 192.168.10.45
```

LISTING 24-8 Remove-IPBlockListEntry cmdlet syntax and usage

Syntax

```
Remove-IPBlockListEntry -Identity IPListEntryId
```

```
[-Server ServerId]
```

Usage

```
Get-IPBlockListEntry | Where {$_.IPRange -eq '192.168.10.45'} |
```

```
Remove-IPBlockListEntry
```

```
Get-IPBlockListEntry | Where {$_.IPRange -eq '192.168.10.0/24'} |
```

```
Remove-IPBlockListEntry
```

Preventing Internal Servers from Being Filtered

Typically, you don't want Exchange to apply Sender ID, content or connection filters to servers on your organization's network or to internal SMTP servers deployed in a perimeter zone. One way to ensure this is to configure message delivery options for your organization's transport servers so that they don't apply filters to IP addresses from internal servers and your perimeter network.

In Exchange Management Shell, you prevent internal servers from being filtered by the Sender ID, content or connection filters by using the `-InternalMailEnabled` parameter of `Set-SenderIdConfig`, `Set-ContentFilterConfig`, and `Set-IPBlockListProvider`. By default, `-InternalMailEnabled` is set to `$false` for these cmdlets, which means authenticated internal email messages aren't processed by the Sender ID filter, the content filter or the connection filter.

Chapter 25. Optimizing Web and Mobile Access

As you'll learn in this chapter, managing web and mobile access is a bit different from other tasks you'll perform as an Exchange administrator—and not only because you use the Microsoft Internet Information Services (IIS) Manager snap-in to perform many of the management tasks. In Exchange 2016, all client connections are handled by Mailbox servers. Mailbox servers provide front-end authentication and proxying, and also perform the actual back-end processing. Generally speaking, clients connect to front-end services and then are routed via local or remote proxy to the back-end endpoint on the Mailbox server that hosts the active copy of the mailbox database storing the user's mailbox.

When users access Exchange mail and public folders over the Internet or a wireless network, virtual directories and web applications hosted by Mailbox servers are working behind the scenes to grant access and transfer files. As you know from previous chapters, Outlook Web App (OWA) lets users access Exchange over the Internet or over a wireless network by using a standard web browser. Exchange ActiveSync lets users access Exchange through a wireless carrier using mobile devices, such as smart phones. Finally, MAPI over HTTP and RPC over HTTP let users access Exchange mailboxes using Microsoft Office Outlook from the Internet.

Navigating IIS Essentials for Exchange Server

On each Mailbox server there is a single instance of IIS that handles front-end and back-end processes. This IIS instance has a Default Web Site with a single virtual directory for each client protocol handled by the server and a corresponding Exchange Back End website with a single virtual directory for each client protocol handled by the server.

IIS handles incoming requests to a website within the context of a web application. A web application is a software program that delivers content to users over HTTP or HTTPS. Each website has a default web application and one or more additional web applications associated with it. The default web application handles incoming requests that you haven't assigned to other web applications. Additional web applications handle incoming requests that specifically reference the application.

Understanding Mobile Access via IIS

When you install an Exchange server, virtual directories and web applications are installed to support various Exchange services. Each web application must have a root virtual directory associated with it. The root virtual directory sets the application's name and maps the application to the physical directory that contains the application's content. Typically, the default web application is associated with the root virtual directory of the website and any additional virtual directories you've created but haven't mapped to other applications.

In the default configuration, the default application handles an incoming request for the / directory of a website as well as other named virtual directories. IIS maps references to / and other virtual directories to the physical directory that contains the related content. For the / directory of the default website, the default physical directory is %SystemRoot%/inetpub/wwwroot.

In most cases, you only need to open port 443 on your organization's firewall to allow users to access Exchange data hosted by IIS. Then you simply tell users the URL that they need to type into their browser's Address field or in their smart phone's browser. Users can then access Outlook Web App or Exchange ActiveSync when they're off-site. The URLs for Outlook Web App and Exchange ActiveSync are different. The Outlook Web App URL is *https://yourserver.yourdomain.com/owa* , and the Exchange ActiveSync URL is *https://yourserver.yourdomain.com/Microsoft-Server-ActiveSync* . Generally, however, the address users enter for both matches the OWA address.

As you configure web and mobile access, don't forget that the infrastructure has two layers:

- A front end that you can customize to control the way users access and work with related services and features
- A back end that handles the back-end processing but that you only modify to control the options that the front end uses for working with the back-end processes

Thus, although you typically modify front-end virtual directories to customize the environment for users, you rarely modify the back-end virtual directories. For example, when you first install Exchange services, Outlook Web App, Exchange Admin Center, and other essential services can only be accessed by clients on the internal network. To allow external clients to access these services, you must specify an external access URL for Outlook Web App, Exchange Admin Center, and other essential services.

Maintaining Virtual Directories and Web Applications

When you install an Exchange server, Exchange Setup installs and configures virtual directories and Web applications for use. The virtual directories and web applications allow authenticated users to access their messaging data from the web. On Mailbox servers, you'll find a Default Web Site that provides front-end services and an Exchange Back End website that provides back-end services. Apps on the front end have corresponding back-end apps, with connections being proxied from the front end to the back end for processing.

In Exchange Management Shell, you can use the `Get-OWAVirtualDirectory` cmdlet to view information about OWA virtual directories, the `New-OWAVirtualDirectory` cmdlet to create an OWA directory if one does not exist, the `Remove-OWAVirtualDirectory` cmdlet to remove an OWA directory, and the `Test-OWAConnectivity` cmdlet to test OWA connectivity. There are similar sets of commands for ActiveSync, Autodiscover, ECP, OAB, Windows PowerShell, MAPI and web services.

Generally, any time you are working with Exchange via the shell, you should specify whether you want to work with the front-end virtual directory or backend virtual directory to ensure the directory you expect to be configured is the one created, modified or removed. Keep the following in mind:

- When you are working with related `Get` commands, only information about front-end directories is provided by default. To also display information about back-end directories, you must use the `-ShowMailboxVirtualDirectories` parameter.
- When you are working with the related `New` and `Remove` commands, you must use the `-Role` parameter to specify whether you are working with front-end or back-end directories. Set `-Role` to `ClientAccess` when you want to configure the front-end virtual directory. Set `-Role` to `Mailbox` when you want to configure the back-end virtual directory.
- With related `Set` commands, you must always specify the identity of the virtual directory that you want to work with. The identity points to the exact endpoint on the front-end or backend.

If you examine the virtual directory structure for the Default Web site or the Exchange Back End website, you'll find several important virtual directories and web applications, including:

- **Autodiscover** Autodiscover is used to provide the Autodiscover service for all clients. By default, this directory is configured for pass-through authentication and the

related app runs within the context of MExchangeAutodiscoverAppPool. For troubleshooting non-configuration issues, use the Autodiscover, Autodiscover.Proxy, and Autodiscover.Protocol health sets. Check these health sets using:

```
Get-ServerHealth ServerId | ?{$_.HealthSetName -match "Autodiscover"} @techjob
```

- **ECP** The Exchange Admin Center (ECP) is used for web-based administration of Exchange. By default, this directory is configured for pass-through authentication and the related app runs within the context of MExchangeECPAppPool. For troubleshooting non-configuration issues, use the ECP and ECP.Proxy and OWA.Protocol health sets. Check the ECP health sets using:

```
Get-ServerHealth ServerId | ?{$_.HealthSetName -match "ECP"}
```

- **EWS** Exchange Web Services (EWS) is used to enable applications to interact with Exchange mailboxes and messaging items using HTTPS. By default, this directory is configured for pass-through authentication and the related app runs within the context of MExchangeServicesAppPool. For troubleshooting non-configuration issues, use the EWS, EWS.Proxy, and EWS.Protocol health sets. Check these health sets using:

```
Get-ServerHealth ServerId | ?{$_.HealthSetName -match "EWS"}
```

- **MAPI** Mapi is the directory that provides MAPI over HTTP services to clients. By default, this directory is configured for pass-through authentication and the related app runs within the context of MExchangeMapiFrontEndAppPool. For troubleshooting non-configuration issues, use the OutlookMapiHttp-related health sets. Check these health sets using:

```
Get-ServerHealth ServerId | ?{$_.HealthSetName -match "MAPI"}
```

- **Microsoft-Server-ActiveSync** Microsoft-Server-ActiveSync is the directory to which Exchange ActiveSync users connect to access their Exchange data. By default, this directory is configured for pass-through authentication and the related app runs within the context of MExchangeSyncAppPool. For troubleshooting non-configuration issues, use the ActiveSync, ActiveSync.Proxy and ActiveSync.Protocol health sets. Check these health sets using:

```
Get-ServerHealth ServerId | ?{$_.HealthSetName -match "ActiveSync"}
```

- **OAB** OAB is the directory that provides the offline address book (OAB) to clients. By default, this directory is configured for pass-through authentication and the related app runs within the context of MExchangeOABAppPool. For troubleshooting non-configuration issues, use the OAB and OAB.Proxy health sets. Check these health sets using:

```
Get-ServerHealth ServerId | ?{$_.HealthSetName -match "OAB"}
```

- **OWA** OWA is the directory to which users connect with their web browsers to start an Outlook Web App session. By default, this directory is configured for pass-through authentication and the related app runs within the context of MExchangeOWAAppPool. For troubleshooting non-configuration issues, use the OWA, OWA.Proxy, and OWA.Protocol health sets. Check these health sets using:

```
Get-ServerHealth ServerId | ?{$_.HealthSetName -match "OWA"}
```

- **PowerShell** PowerShell is the directory to which the Exchange Management tools connect for remote administration. Depending on whether you are working with the front-end or back-end, the related app runs within the context of `MSExchangePowerShellFrontEndAppPool` or `MSExchangePowerShellBackEndAppPool`. For troubleshooting non-configuration issues, use the `RPS` and `RPS.Proxy` health sets. Check these health sets using:

```
Get-ServerHealth ServerId | ?{$_.HealthSetName -match "RPS"}
```

- **RPC** RPC is the directory that provides Remote Procedure Call (RPC) services to clients. By default, this directory is configured for pass-through authentication and the related app runs within the context of `MSExchangeRPCProxyAppPool`. Whether the `Rpc` virtual directory on the front end connects to the `Rpc` virtual directory or the `RpcWithCert` virtual directory on the backend depends on whether an SSL certificate is used as part of authentication.
- **Public** Public is the directory to which mailbox users are connected to access the default Public Folders tree. This directory exists only on Mailbox servers and doesn't have a specifically configured application pool. For troubleshooting non-configuration issues, use the `PublicFolders` and `OWA` health sets. Check the `PublicFolders` health set using:

```
Get-ServerHealth ServerId | ?{$_.HealthSetName -eq "PublicFolders"}
```

For troubleshooting configuration issues with virtual directories, you might need to remove and recreate the front-end virtual directory first, and then check to see if this resolves the problem before removing and recreating the back-end virtual directory. As an example, if you've determined the OWA virtual directory is misconfigured, you can remove it by using `Remove-OwaVirtualDirectory`, and then recreate it by using `New-OwaVirtualDirectory`. You could remove and then recreate the OWA virtual directory from the Default Web Site on MailServer21 by using the following commands:

```
remove-owavirtualdirectory -identity "mailserver21\owa (Default Web Site)"
```

```
new-owavirtualdirectory -server mailserver 21  
-websitename "Default Web Site"
```

By default, the `New-OwaVirtualDirectory` and `New-EcpVirtualDirectory` commands enable basic authentication and forms authentication but do not enable Windows authentication. As Windows authentication is required for OWA and ECP, you'll want to use the `Set-OwaVirtualDirectory` and `Set-EcpVirtualDirectory` commands to modify the default authentication settings. In the following example, you enable Windows authentication and disable basic and forms authentication:

```
set-owavirtualdirectory -identity "mailserver21\owa (Default Web Site)"  
-WindowsAuthentication $True -Basicauthentication $false  
-Formsauthentication $false
```

TIP You can set properties on some or all virtual directories by piping the

output of `Get-OwaVirtualDirectory` to `Set-OwaVirtualDirectory`. For example, the following command allows users to change their passwords by default for all Outlook Web App virtual directories:

```
Get-OwaVirtualDirectory | Set-OwaVirtualDirectory -ChangePassword $true
```

After you recreate a virtual directory you should restart IIS services. You can do this in IIS Manager or by entering the following command at an elevated command prompt or shell:

```
iisreset
```

You can diagnose non-configuration problems with a particular feature as well as any related proxy and protocol features by using `Get-HealthReport` to check the status of the related health sets. Try to resolve the problem by using the following techniques while verifying the issue still exists each time you take a corrective action:

1. Try to isolate the problem to a specific server by running a health check for the feature on each Exchange server. If you find an Unhealthy status for the feature on a particular server, you've likely isolated the problem and identified the server experiencing the problem and can skip Steps 2 and 3.
2. If you are unable to isolate the problem to a specific server or servers, try to log on to OWA or ECP, and then use the feature using the URL for a specific Mailbox server. If this fails, try accessing and log on to a different Mailbox server. This should help you verify whether the problem is with a particular Mailbox server. Remember that the Mailbox server used is the one that contains the mailbox database where the mailbox for the user is stored.
3. Using the Services console, verify that all essential Exchange services are running on the Mailbox servers. If an essential service isn't running, select it and then click Start.
4. Verify network connectivity between the Mailbox servers. One way to do this is to log on to each server and try to ping the other servers. If you correct a connectivity issue, check to see if the issue is resolved. Most features require connectivity to domain controllers.
5. In IIS Manager, connect to the server that's reporting the health issue or otherwise experiencing a problem with the feature you are troubleshooting. Expand the Sites node and verify that the Default Web Site or Exchange Back End website is running as appropriate. If a required website isn't running, click Start in the Actions pane to start it. This should resolve the problem.
6. Under Application Pools, verify that the required application pools have been started. If a required application pool hasn't been started, select it and then click Start in the Actions pane.
7. If you suspect an issue with a required application pool on the front-end server, the back-end server, or both, select the application pool and then click Recycle in the Actions pane to recycle its work processes.

8. If the problem isn't resolved yet, restart the website in which the problem is occurring or the IIS itself. To restart a website, select the website in IIS Manager, and then select Restart in the Actions pane. To restart IIS, select the server node in IIS Manager, and then Restart in the Actions pane.
9. If the problem still isn't resolved, restart the server. If restarting the server doesn't resolve the problem, you likely have a configuration problem that can be resolved by removing and recreating the related virtual directories.

Starting, Stopping, and Restarting Websites

Websites run under a server process that you can start, stop, and pause, much like other server processes. For example, if you're changing the configuration of a website or performing other maintenance tasks, you might need to stop the website, make the changes, and then restart it. When a website is stopped, it doesn't accept connections from users and can't be used to deliver or retrieve mail.

The master process for all websites is the World Wide Web Publishing Service. Stopping this service stops all websites using the process, and all connections are disconnected immediately. Starting this service restarts all websites that were running when you stopped the World Wide Web Publishing Service.

You can start, stop, or restart a website by completing the following steps:

1. Start IIS Manager. In Server Manager, click Tools, and then select Internet Information Services (IIS) Manager.



2. In IIS Manager, expand the server and related Sites node by double-clicking the entry for the server you want to work with, and then double-clicking Sites.
3. Select the website you want to manage. Using the options in the Actions pane, you can now do the following:
 - Select Start to start the website.
 - Select Stop to stop the website.
 - Select Restart to stop and then start the website.

If you suspect there's a problem with the World Wide Web Publishing Service or other related IIS services, you can use the following technique to restart all IIS services:

1. Start IIS Manager. In Server Manager, click Tools, and then select Internet Information Services (IIS) Manager.
2. Select the entry for the server you want to work with, and then select Restart in the Actions Pane.



Configuring Outlook Web App Features

Microsoft uses the term *segmentation* to refer to your ability to enable and disable the various features within Outlook Web App. Segmentation settings applied to the OWA virtual directory control the features available to users. If a server has multiple OWA virtual directories or you have multiple Mailbox servers, you must configure each directory and server separately.

Managing Segmentation Features

A summary of the segmentation features that are enabled by default for use with Outlook Web App follows:

- **All Address Lists** When enabled, users can view all the available address lists. When this feature is disabled, users can view only the default global address list.
- **Calendar** When enabled, users can access their calendars in Outlook Web App.
- **Change Password** When enabled, users can change their passwords in Outlook Web App.
- **Contacts** When enabled, users can access their contacts in Outlook Web App.
- **Direct File Access** When enabled, users can open attachments directly.
- **Email Signature** When enabled, users can customize their signatures and include a signature in outgoing messages.
- **Exchange ActiveSync** When enabled, users can remove mobile devices, initiate mobile wipe, view their device passwords, and review their mobile access logs.
- **Inbox Rules** When enabled, users can customize rules in Outlook Web App.
- **Instant Messaging** When enabled, users can access Instant Messaging in Outlook Web App.
- **Journaling** When enabled, the Journal folder is visible in Outlook Web App.
- **Junk Email Filtering** When enabled, users can filter junk email using Outlook Web App.
- **Notes** When enabled, users can access their notes in Outlook Web App.
- **Premium Client** When enabled, users can use the standard version of Outlook Web App (or the light version if that is the version supported by their browser). When this feature is disabled, users can only access the light version of Outlook Web App.
- **Public Folders** When enabled, users can browse and read items in public folders using Outlook Web App.
- **Recover Deleted Items** When enabled, users can view items that have been deleted from Deleted Items and choose whether to recover them.
- **Reminders And Notifications** When enabled, users can receive new email notifications, task reminders, calendar reminders, and automatic folder updates.
- **Tasks** When enabled, users can access their tasks in Outlook Web App.
- **Text Messaging** When enabled, users can send and receive text messages and create text message notifications in Outlook Web App.

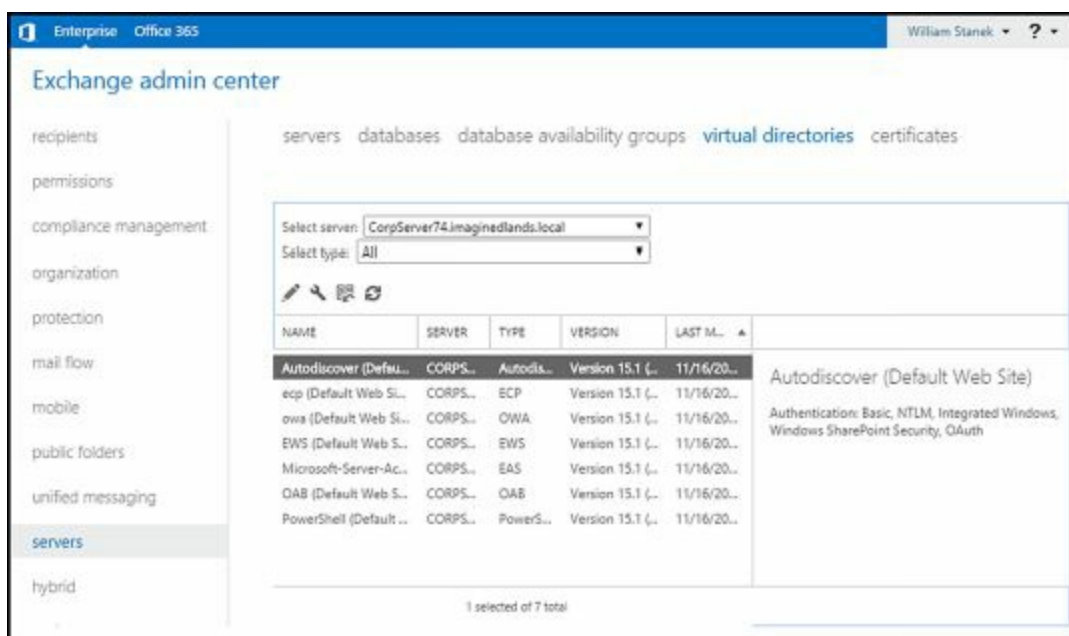
- **Themes** When enabled, users can change the color scheme in Outlook Web App.
- **Unified Messaging** When enabled, users can access their voice mail and faxes in Outlook Web App. They can also configure voice mail options.
- **Office Web Apps Viewing** When enabled, users can view supported file types while using OWA in their web browser.


You manage segmentation features in several ways:

- In Exchange Management Shell, you can enable or disable segmentation features on a per server basis by running the Set-OWAVirtualDirectory cmdlet on Mailbox servers.
- In Exchange Admin Center and Exchange Management Shell, you can define Outlook Web App policies that enable or disable segmentation features and then apply these policies to users. Settings in Outlook Web App policies override virtual directory settings.
- In Exchange Management Shell, you can enable or disable segmentation features for individual users by using the Set-CASMailbox cmdlet. These settings override settings applied through policies and virtual directories.

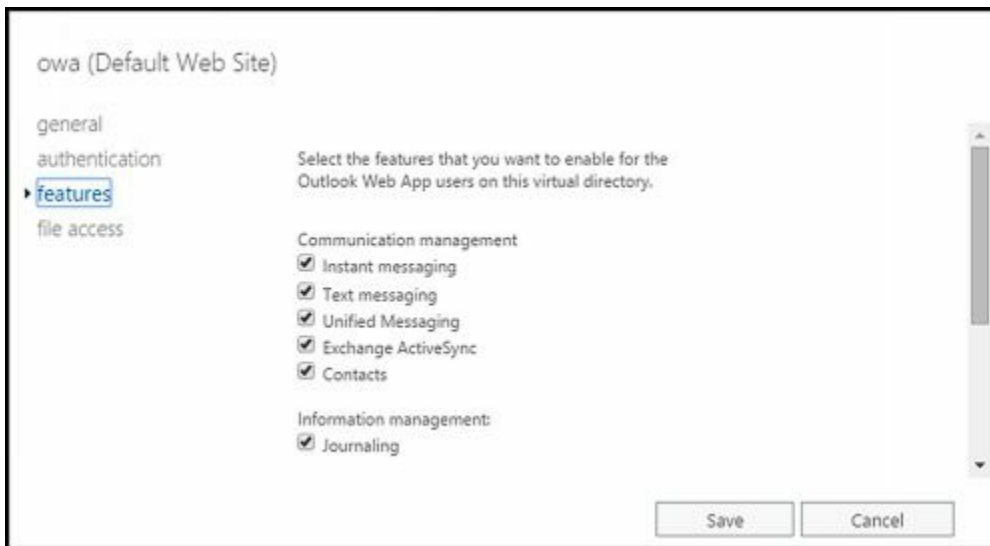
To enable or disable segmentation features for a particular virtual directory, complete the following steps:

1. In the Exchange Admin Center, select Servers in the Navigation menu, and then select Virtual Directories to view a list of the front-end virtual directories used by all Exchange servers in the organization. To streamline the view, select a specific server on the Select Server list and/or choose a specific directory type on the Select Type list.



2. Select the OWA virtual directory you want to configure, and then select Edit ().

3. In the Virtual Directory dialog box, select the Features page. Scroll down and then click More Options to view all the features you can manage.

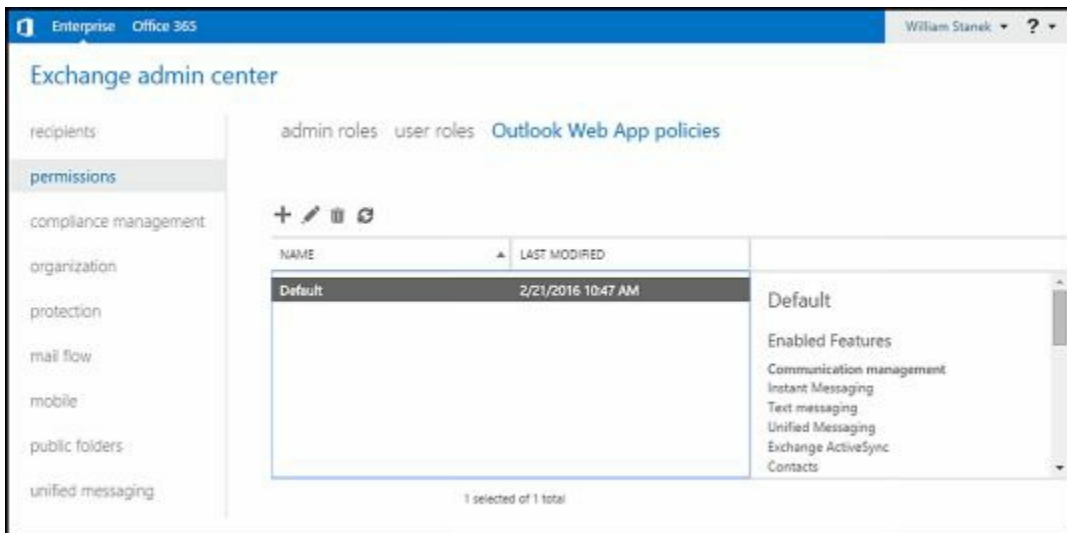


4. By default, all features are enabled. To disable a feature, clear the related checkbox.
5. By default users can view web-ready documents in their browser and open attachments directly whether they are using a public or private computer. As necessary, use the options on the File Access page to change the file access options.
6. Click Save to apply the settings.




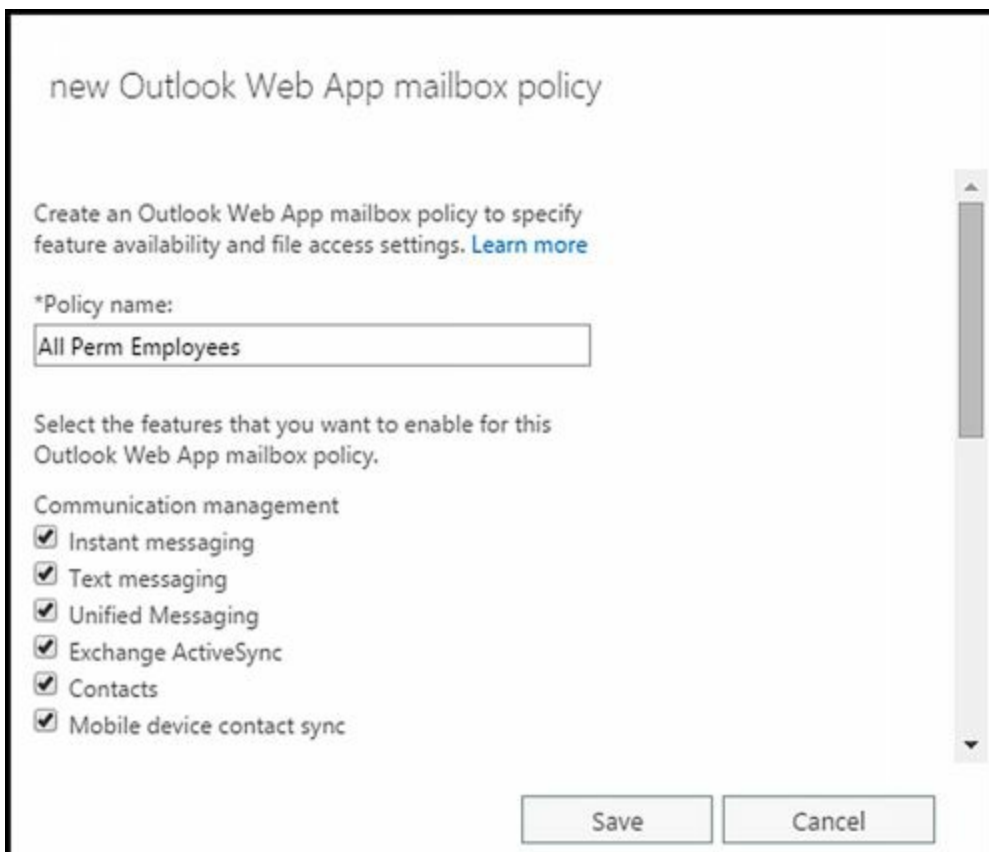
Managing Outlook Web App Policies

In the Exchange Admin Center, select Permissions in the Navigation menu, and then select Outlook Web App Policies to view the currently defined policies. Select a policy to view its settings in the details pane.



To create an Outlook Web App policy, follow these steps:

1. When you select Permissions > Outlook Web App Policies in Exchange Admin Center, you'll see a list of current policies. You can create a new policy by clicking Add ().
2. In the Policy Name text box, type a descriptive name for the policy, such as All Permanent Employees.
3. To view all of the features you can configure, click More Options.
4. By default, all features are enabled. To disable a feature, clear the related checkbox.
5. Click Save to create the policy.



In Exchange Management Shell, you can create Outlook Web App policies by using `New-OwaMailboxPolicy` and then set the properties of the policy by using `Set-OwaMailboxPolicy`. The following example creates a policy called `AllUsers` and then configures its settings:

```
New-OwaMailboxPolicy -Name AllUsers
```

```
Set-OwaMailboxPolicy -Identity AllUsers -AllAddressLists $false  
-ChangePasswordEnabled $false -AllowOfflineOn "NoComputers"  
-ContactsEnabled $false -FacebookEnabled $false -LinkedInEnabled $false  
-CalendarEnabled $true
```

Use `Get-OwaMailboxPolicy` to confirm that the properties of the policy are set as expected. Afterward, you can apply the policy to users by using the `-OwaMailboxPolicy` property of `Set-CASMailbox`. Listing 25-1 shows various ways you can apply the policy.

LISTING 25-1 Techniques for applying OWA mailbox policies to mailbox users

Apply the policy to the mailbox user named HenryJ

```
Set-CASMailboxPolicy -Identity HenryJ -OwaMailboxPolicy "AllUsers"
```

Apply the policy to every mailbox in the Exchange organization

```
Get-Mailbox -ResultSize Unlimited | Set-CASMailboxPolicy  
-OwaMailboxPolicy "AllUsers"
```

Apply the policy to every mailbox in the Sales database

```
Get-MailboxDatabase "Sales" | Get-Mailbox -ResultSize Unlimited |  
Set-CASMailboxPolicy -OwaMailboxPolicy "AllUsers"
```

Apply the policy to all mailboxes in every mailbox database on MailboxServer18

```
Get-Mailbox -Server MailboxServer18 -ResultSize Unlimited  
Set-CASMailboxPolicy -OwaMailboxPolicy "AllUsers"
```

Managing Bindings, Connections and Authentication

As you optimize IIS for your Exchange environment, you'll want to look at the configuration settings for bindings, SSL and connections. You might also need to redirect users to alternate URLs, control access and modify authentication settings. These topics are discussed in this section.

Optimizing the Mobile Access Websites

Each website hosted by IIS has one or more bindings. A binding is a unique combination of ports, IP addresses, and host names that identifies a website. For unsecure connections, the default port is TCP port 80. For secure connections, the default port is TCP port 443. The default IP address setting is to use any available IP address. The default host name is the Mailbox server's DNS name.

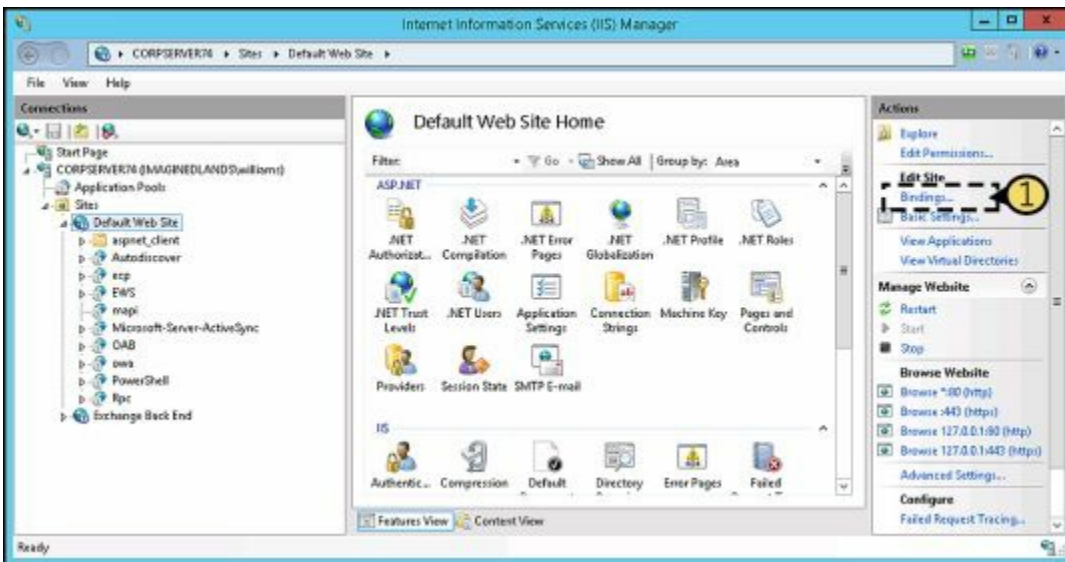
Normally, you wouldn't want to multihome a Mailbox server; however, when the server is multihomed, or when you use it to provide Outlook Web App or Exchange ActiveSync services for multiple domains, the default configuration isn't ideal. On a multihomed server, you'll usually want messaging protocols to respond only on a specific IP address. To do this, you need to change the default settings. On a server that provides Outlook Web App and Exchange ActiveSync services for multiple domains, you'll usually want to specify an additional host name for each domain.

When you are working with IIS, you can change the identity of a website by completing the following steps:

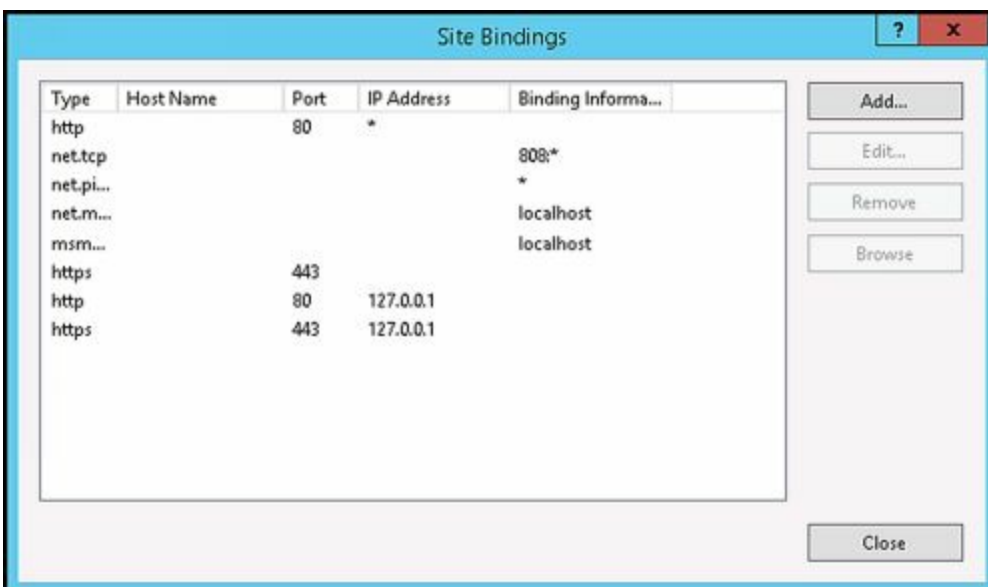
1. If you want the website to use a new IP address, you must configure the IP address on the server before trying to specify it on the website.
2. Start IIS Manager. In Server Manager, click Tools, and then select Internet Information Services (IIS) Manager.

NOTE By default, IIS Manager connects to the services running on the local computer. If you want to connect to a different server, select the Start Page node in the left pane, and then click the Connect to a Server link to start the Connect To Server Wizard. Follow the prompts to connect to the remote server.

3. In IIS Manager, expand the server and related Sites node by double-clicking the entry for the server with which you want to work, and then double-clicking Sites.



- In the left pane, select the website that you want to manage, and then select Bindings on the Actions pane. You can now use the Site Bindings dialog box to configure multiple bindings for the website.



- Use the Site Bindings dialog box to manage the site's bindings by using the following settings:

- Add** Adds a new identity. To add a new identity, click Add. In the Add Site Bindings dialog box, select the binding type, IP address, and TCP port to use. Optionally, type a host header name or select a Secure Sockets Layer (SSL) certificate as appropriate for the binding type. Click OK when you have finished.
- Edit** Allows you to edit the currently selected identity. To edit an identity, click the identity, and then click Edit. In the Edit Site Bindings dialog box, select an IP address and TCP port to use. Optionally, type a host header name or select an SSL certificate as appropriate for the binding type. Click OK when you have finished.
- Remove** Allows you to remove the currently selected identity. To remove an identity, click the identity, and then click Remove. When prompted to confirm, click Yes.
- Browse** Allows you to test an identity. To test an identity, click the identity, and then click Browse. IIS Manager then opens a browser window and connects to the

selected binding.

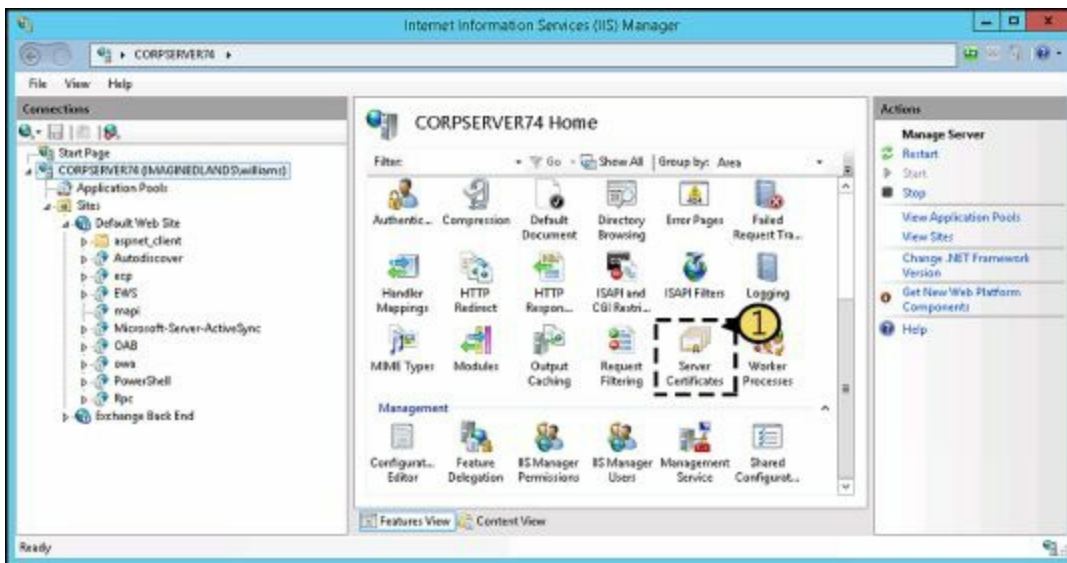
6. Click Close.

Enabling SSL on Websites

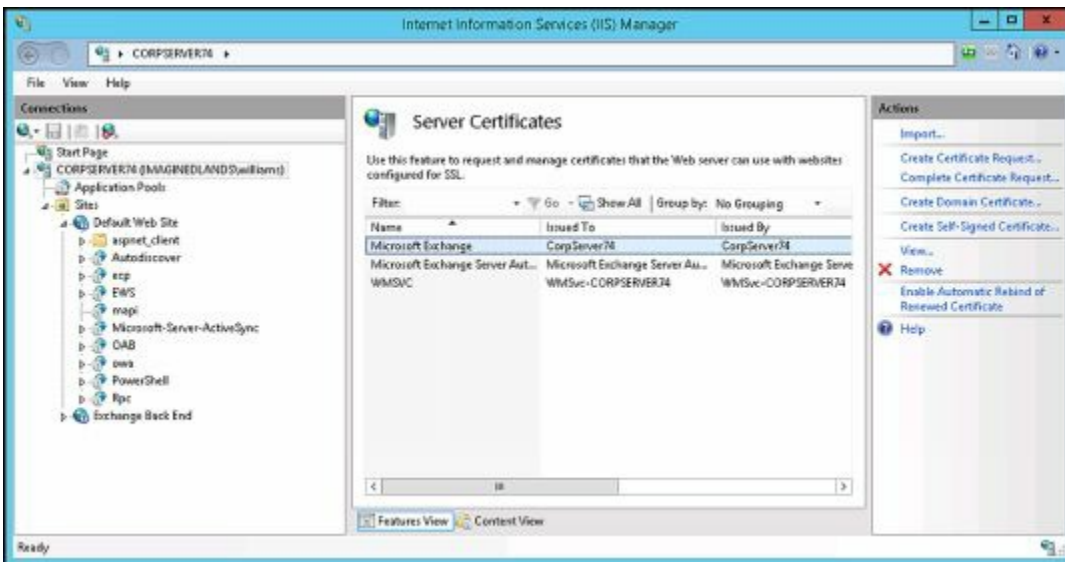
SSL is a protocol for encrypting data that is transferred between a client and a server. Without SSL, servers pass data in readable, unencrypted text to clients, which could be a security risk in an enterprise environment. With SSL, servers pass data encoded using encryption.

Although websites are configured to use SSL on port 443 automatically, the server won't use SSL unless you've created and installed a valid X.509 certificate. When you install an Exchange server, a default X.509 certificate is created for Exchange Server 2016 and registered with IIS. In IIS Manager, you can view the default X.509 certificate by completing the following steps:

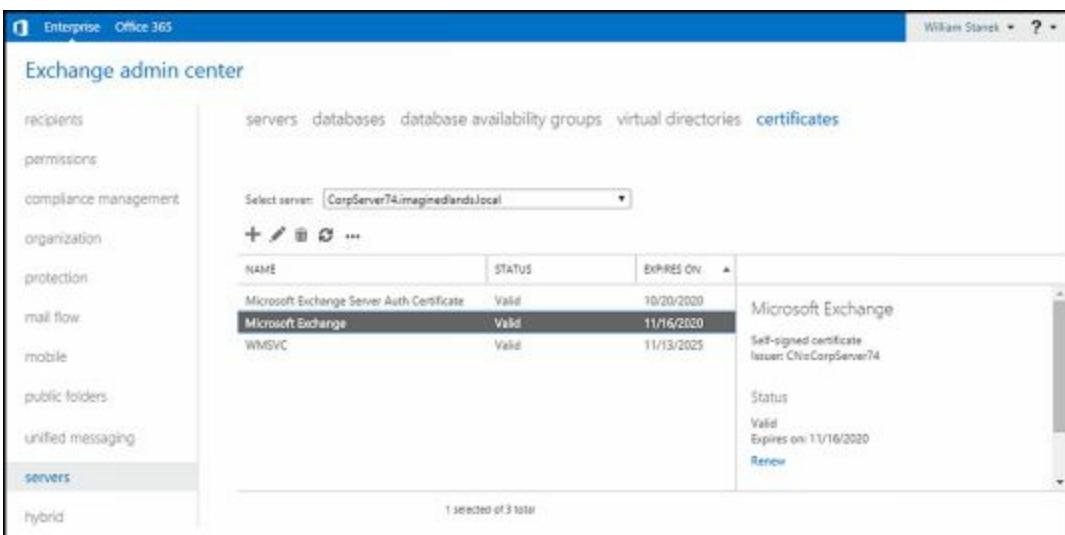
1. Log on to the Mailbox server. Start IIS Manager. In Server Manager, click Tools, and then select Internet Information Services (IIS) Manager.
2. In IIS Manager, expand the Default server node, and then double-click the Server Certificates feature.



3. On the Server Certificates page, you'll see a list of certificates the web server can use. The default X.509 certificate for Exchange Server has the name Microsoft Exchange. Click the certificate entry, and then click View in the Actions pane to view detailed information regarding the certificate. By default, this certificate is valid for one year from the date you install the server.




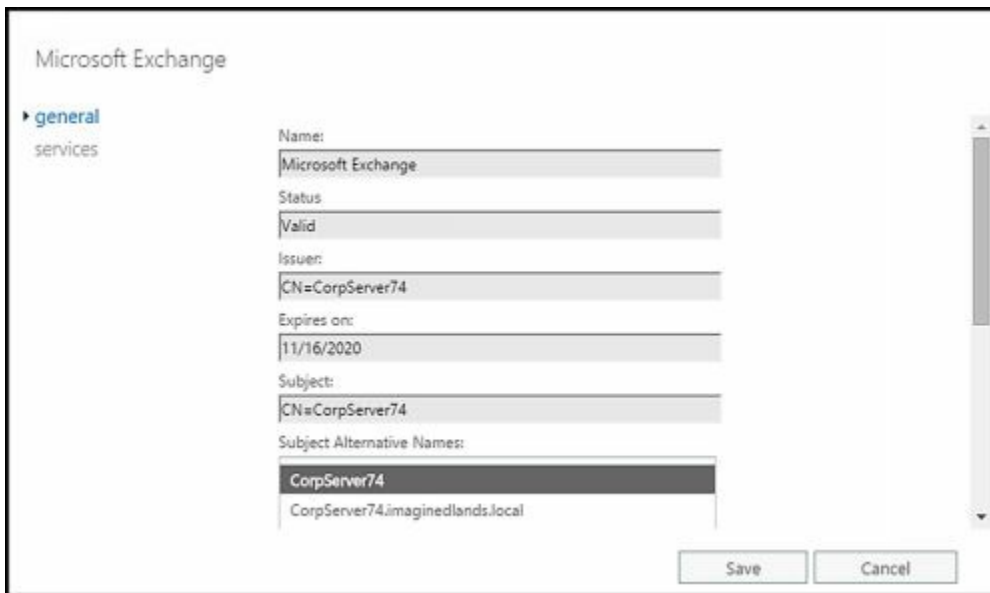
For a long-term solution, you need to create a permanent certificate for the server. This certificate can be a certificate assigned by your organization’s certificate authority (CA) or a third-party certificate.



To create a certificate for use with Exchange and IIS, use the features provided by the Exchange management tools. In the Exchange Admin Center, you can view available certificates for Exchange servers by selecting Servers in the Navigation menu, and then selecting Certificates. Next, on the Select Server list, choose the server you want to work with. You’ll then see a list of available certificates for this server.

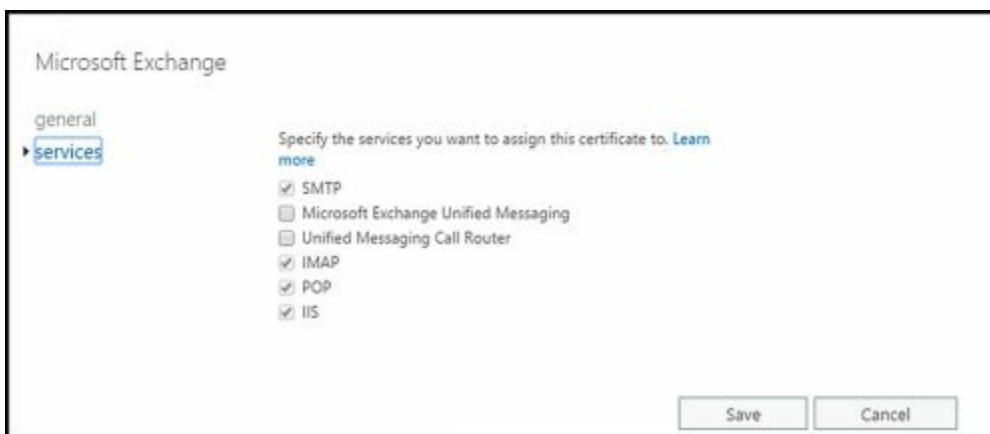
You can view the general settings for the certificate by selecting it and then selecting

Edit (). The subject alternative names associated with the certificate determines the names that can be used when establishing SSL connections. Typically, the subject alternative names include the host name and the fully-qualified domain name of the server.



CAUTION Don't make any changes to certificates as this could invalidate them. When you are finished viewing a certificate, click Cancel to exit the properties dialog box without saving any changes. The default certificates were created by using Exchange Management Shell and should only be modified or renewed by using Exchange Management Shell. The same is true for any other certificate created using the shell.

On the Services page, each selected option represents a service assigned to the certificate. By assigning a service to a certificate, you are allowing the certificate to be used to secure the service. After you are done viewing a certificate's properties, click Cancel (you don't want to inadvertently make any changes to a certificate).



To request and create a certificate from a certification authority, complete the following steps:

1. In the Exchange Admin Center, select Servers in the Navigation menu, and then select Certificates.
2. On the Select Server list, choose the server with which you want to work.
3. Click Add to start the New Exchange Certificate Wizard.
4. Select Create A Request to use the wizard to create a certificate request file, and then click Next.

new Exchange certificate

This wizard will create a new certificate or a certificate request file.
You can either create a self-signed certificate or request a certificate from a certification authority. [Learn more--](#)

Create a request for a certificate from a certification authority
 Create a self-signed certificate

Next Cancel

5. Type a descriptive name for the certificate, and then click Next.

new Exchange certificate

*Friendly name for this certificate:

Exchange Cert from CA

Back Next Cancel

6. Specify the root domain for the certificate. If you want the certificate to be usable for all subdomains of your root domain, select the Request A Wildcard Certificate checkbox, and then click Next.

new Exchange certificate

Request a wildcard certificate. A wildcard certificate can be used to secure all subdomains under your root domain with a single certificate. [Learn more](#)

*Root domain:

imaginedlands.com

Back Next Cancel

7. Click Browse. Choose the server where you want to store the request. Typically, this is the server where you will install the certificate. Click Next.

new Exchange certificate

*Store certificate request on this server:

CORPSEVER74

- Identify your organization by entering the organization name, department name, city, state, and country. These values are all required and must be entered before you can continue. Click Next.

new Exchange certificate

Specify information about your organization. This is required by the certification authority. [Learn more](#)

*Organization name:
Imagined Lands Ltd.

*Department name:
Tech Team

*City/Locality:
Seattle

*State/Province:
Washington

*Country/Region name:
United States ▼

- Specify the full file path for a network location where the certificate request file can be saved, such as \\CorpServer74\Data\CertRequest.req. Click Finish.

new Exchange certificate

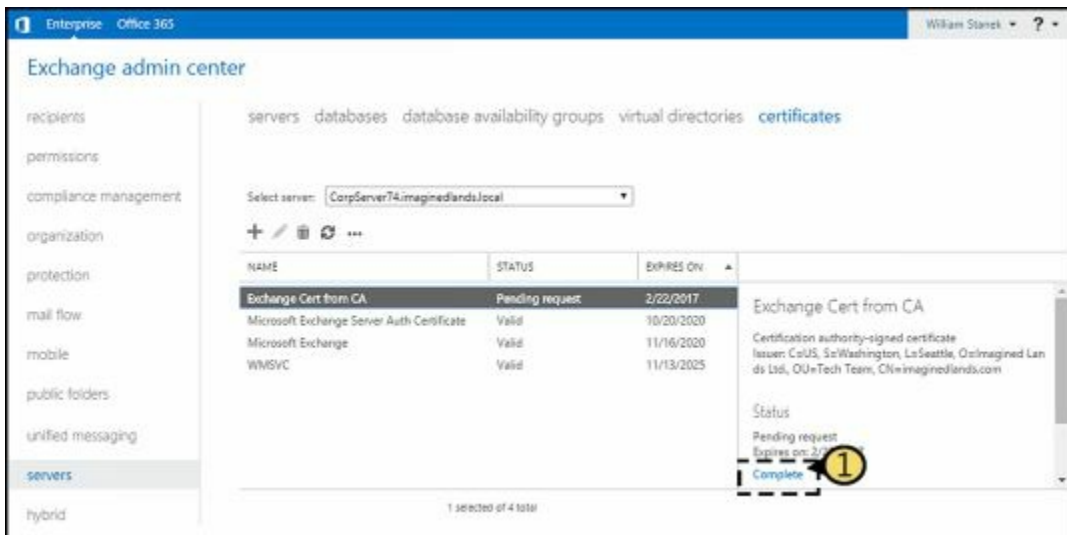
*Save the certificate request to the following file (example:
\\myservername\share\mycertrequest.REQ):

\\CorpServer74\Data\CertRequest.req

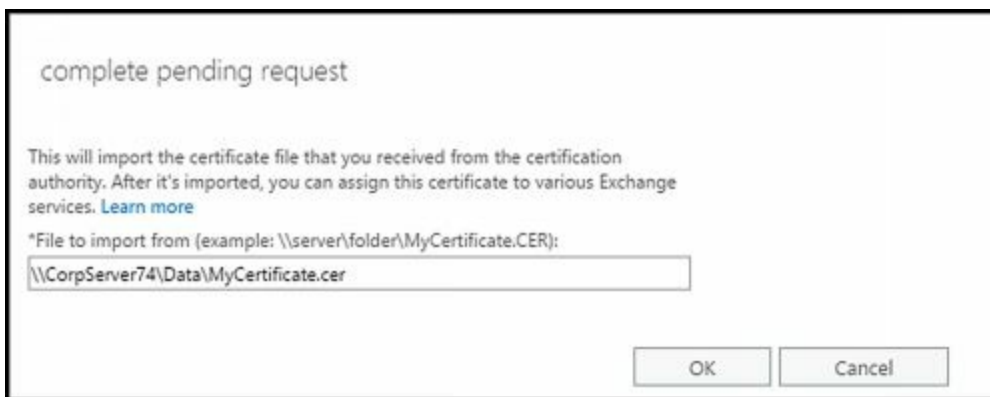
You'll need to submit the contents of the file you entered to a certification authority.

After you receive the certificate file from the certification authority, you'll need to click Complete in the Information pane to install it on your Exchange server. [Learn more](#)

Send the certificate request file to a third-party certificate authority or your organization's CA as appropriate. When you receive the certificate back from the CA, import the certificate. In the Certificates area, you'll see an entry for the certificate with a status of Pending Request. Select this entry, and then select Complete in the details pane.

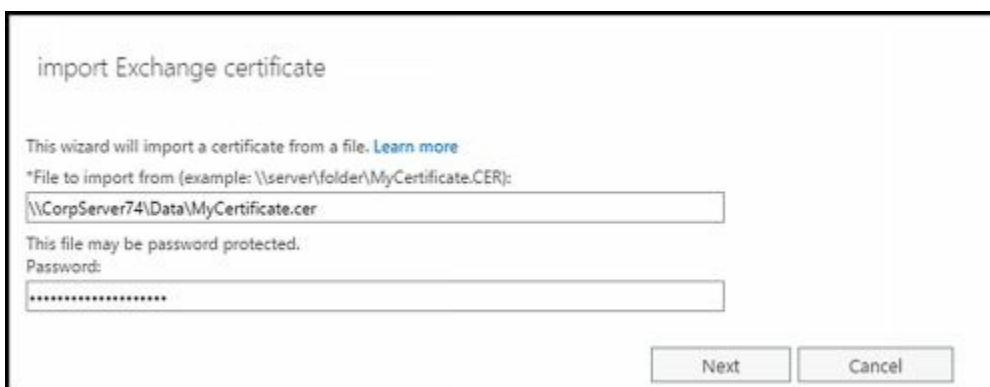


Next, in the Complete Pending Request dialog box, specify the full file path for a network location where the certificate file is available to be imported, such as \\CorpServer74\Data\MyCertificate.cer. Click OK.




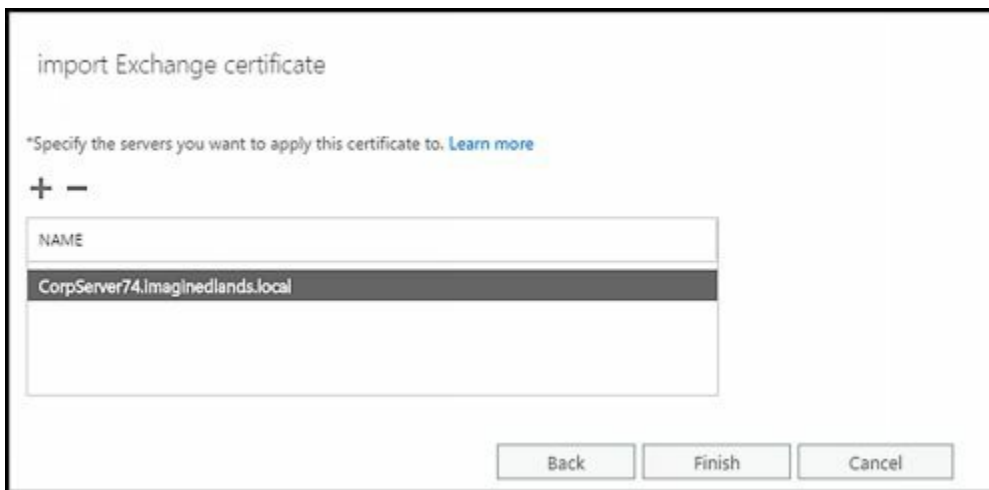
If you have a certificate to install but don't have a pending request, you can import the certificate while working with the Certificates area in the Exchange Admin Center as well. To do this, complete the following steps:

1. Click the More button (**...**), and then select Import Exchange Certificate to start the Import Exchange Certificate Wizard. Use the wizard to import the certificate file.



2. Specify the full file path for a network location where the certificate file is available to be imported, such as \\CorpServer74\Data\MyCertificate.cer. If the file is password-protected, enter the password. Click Next.

3. Click Add (). In the Select A Server dialog box, select a server to which the certificate should be applied, and then click Add. Repeat this process to add additional servers. Click OK.
4. Click Finish to import the certificate.



After you've installed the certificate, you should test the certificate with an external client by accessing OWA from a remote computer. Clients won't automatically trust self-signed certificates or certificates issued by your CA; therefore, you might see an error stating that there is a problem with the website's security certificate. In this case, you'll need to configure the client to trust the certificate. One way to do this with Internet Explorer is follow these steps:

1. Click the Continue To This Website link. When you continue to the site, a Certificate Error option appears to the right of the address field.
2. Click the Certificate Error to display a related error dialog box, and then click View Certificates to display the Certificate dialog box.
3. On the General tab of the Certificate dialog box, you'll see an error stating the CA Root Certificate isn't trusted. Note the certificate details.
4. To enable trust, you must install this certificate in the Trusted Root Certification Authorities store on the computer. The browser will then trust the certificate, and you shouldn't see the certificate error again for this client.



You also can test services supported by the certificate. Test web services by using Test-OutlookWebServices as shown in the following example:

```
test-outlookwebservices | fl
```

By default Test-OutlookWebServices, verifies the Availability service, Outlook Anywhere, Offline Address Book, and Unified Messaging. You can test OWA and ECP by using Test-OwaConnectivity and Test-EcpConnectivity respectively.

Another way to test connectivity is to use the Remote Connectivity Analyzer, which is accessed by entering the following URL in a web browser:

<https://testexchangeconnectivity.com> .

Restricting Incoming Connections

You can control incoming connections to a website in several ways including setting a maximum limit on the bandwidth used, setting a limit on the number of simultaneous connections, and setting a connection time-out value. However, you typically wouldn't want to perform any of these actions for an Exchange server or OWA. OWA has its own timers based on whether the end user is on a public/shared or a private computer. These values are fixed and not affected by any restrictions or settings discussed in this section.

Normally, websites do not have maximum bandwidth limits and accept an unlimited number of connections, which is an optimal setting in most environments. However,

when you're trying to prevent the underlying server hardware from becoming overloaded or you want to ensure other websites on the same computer have enough bandwidth, you might want to limit the bandwidth available to the site and the number of simultaneous connections. When either limit is reached, no other clients are permitted to access the server. The clients must wait until the connection load on the server decreases.

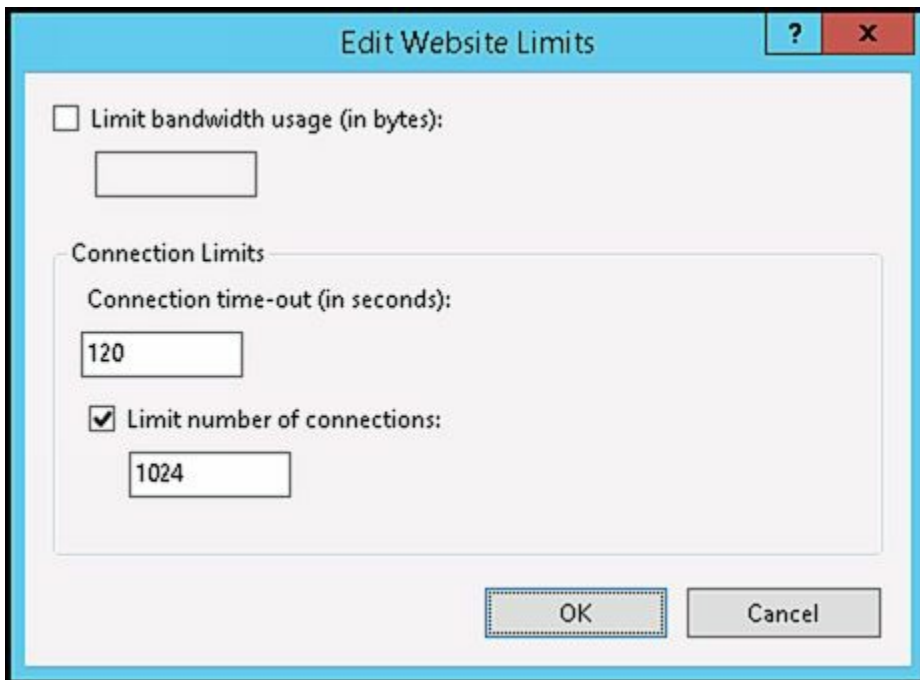
The connection time-out value determines when idle user sessions are disconnected. With the default website, sessions time out after they've been idle for 120 seconds (2 minutes). It's a sound security practice to disconnect idle sessions and force users to log back on to the server. If you don't disconnect idle sessions within a reasonable amount of time, unauthorized persons could gain access to your messaging system through a browser window left unattended on a remote terminal.

You can modify connection limits and time-outs by completing the following steps:

1. Start IIS Manager. In Server Manager, click Tools, and then select Internet Information Services (IIS) Manager.
2. In IIS Manager, expand the server and related Sites node by double-clicking the entry for the server with which you want to work, and then double-clicking Sites.
3. In the left pane, select the website that you want to manage, and then click Limits in the Actions pane. This displays the Edit Website Limits dialog box.



4. To remove maximum bandwidth limits, clear the Limit Bandwidth Usage check box. To set a maximum bandwidth limit, select the Limit Bandwidth Usage check box, and then set the desired limit in bytes.



5. The Connection Time-Out field controls how long idle user sessions remain connected to the server. The default value is 120 seconds. Type a new value to change the current time-out value.
6. To remove connection limits, clear the Limit Number Of Connections check box. To set a connection limit, select the Limit Number Of Connections check box, and then type a limit.
7. Click OK.

Redirecting Users to Alternate Urls

You might occasionally find that you want to redirect users to alternate URLs. For example, you might want users to type **http://mail.tvpress.com** and get redirected to *https://mail.tvpress.com/owa* .

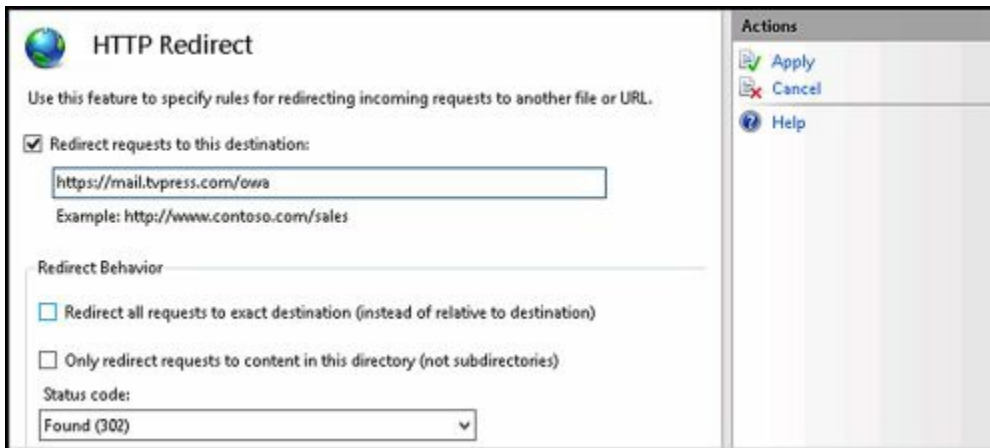
You can redirect users from one URL to another by completing the following steps:

1. Start IIS Manager. In Server Manager, click Tools, and then select Internet Information Services (IIS) Manager.
2. In IIS Manager, navigate to the level you want to manage. You manage redirection for an entire site at the site level, and redirection for a directory at the directory level.
3. In the main pane, double-click the HTTP Redirect feature. This displays the HTTP Redirect page.

NOTE With IIS, HTTP redirection is an optional role service. Therefore, if the HTTP Redirect feature is not available, you need to install the related role service by using Server Manager's Add Roles And Features Wizard.

4. On the HTTP Redirect page, select Redirect Requests To This Destination.
5. In the Redirect Requests To This Destination text box, type the URL to which the user should be redirected. To redirect the user to a different server, type the full

path, starting with `http://` or `https://`, such as `https://mailer2.tvpress.com/owa`. To redirect the user to a virtual directory on the same server, type a slash mark (/) followed by the directory name, such as `/owa`. Click Apply to save your settings.



The screenshot shows the 'HTTP Redirect' configuration window. The title bar reads 'HTTP Redirect'. Below the title, there is a description: 'Use this feature to specify rules for redirecting incoming requests to another file or URL.' The main configuration area includes a checked checkbox 'Redirect requests to this destination:' followed by a text box containing 'https://mail.tvpress.com/owa'. Below this, an example is shown: 'Example: http://www.contoso.com/sales'. Under the 'Redirect Behavior' section, there are two unchecked checkboxes: 'Redirect all requests to exact destination (instead of relative to destination)' and 'Only redirect requests to content in this directory (not subdirectories)'. At the bottom, there is a 'Status code:' label and a dropdown menu currently set to 'Found (302)'. On the right side, there is an 'Actions' pane with three buttons: 'Apply' (with a green checkmark icon), 'Cancel' (with a red X icon), and 'Help' (with a blue question mark icon).

Controlling Access to the HTTP Server

IIS supports several authentication methods, including the following:

- **Anonymous authentication** With anonymous authentication, IIS automatically logs users on with an anonymous or guest account. This allows users to access resources without being prompted for user name and password information.
- **ASP.NET Impersonation** With ASP.NET Impersonation, a managed code application can run either as the user authenticated by IIS or as a designated account that you specify when configuring this mode.
- **Basic authentication** With basic authentication, users are prompted for logon information. When entered, this information is transmitted unencrypted (base64-encoded) across the network. If you've configured secure communications on the server, as described in the section of this chapter titled "Enabling SSL on Websites," you can require that clients use SSL. When you use SSL with basic authentication, the logon information is encrypted before transmission.
- **Digest authentication** With digest authentication, user credentials are transmitted securely between clients and servers. Digest authentication is a feature of HTTP 1.1 and uses a technique that can't be easily intercepted and decrypted.
- **Forms authentication** With Forms authentication, you manage client registration and authentication at the application level instead of relying on the authentication mechanisms in IIS. As the mode name implies, users register and provide their credentials using a logon form. By default, this information is passed as cleartext. To avoid this, you should use SSL encryption for the logon page and other internal application pages.
- **Windows authentication** With Windows authentication, IIS uses kernel-mode Windows security to validate the user's identity. Instead of prompting for a user name and password, clients relay the logon credentials that users supply when they log on to Windows. These credentials are fully encrypted without the need for SSL, and they include the user name and password needed to log on to the network.

When you install IIS on a Mailbox server, you are required to enable basic

authentication, digest authentication, and Windows authentication. These authentication methods, along with anonymous authentication, are used to control access to the server's virtual directories. A virtual directory is simply a folder path that is accessible by a URL. For example, you could create a virtual directory called Data that is physically located on C:\CorpData\Data and accessible by using the URL *https://myserver.mydomain.com/Data* .

The default authentication settings for important virtual directories are as follows:

- [ActiveSync](#) has basic authentication enabled for the front-end and back-end.
- [Autodiscover](#) has anonymous, basic, and Windows authentication enabled for the front-end and back-end.
- [ECP](#) has anonymous and basic authentication enabled for the front-end and back-end.
- [EWS](#) has anonymous and windows authentication enabled for the front-end and back-end.
- [Mapi](#) has Windows authentication enabled on the front-end; anonymous authentication enabled on the back-end.
- [OAB](#) has Windows authentication enabled for the front-end and back-end.
- [OWA](#) has basic authentication enabled on the front-end; anonymous and Windows authentication are enabled on the back-end.
- [PowerShell](#) doesn't have any authentication methods enabled on the front-end; Windows authentication is enabled on the back-end.
- [Rpc](#) has basic and Windows authentication enabled for front-end and back-end.

You should rarely change the default settings. However, if your organization has special needs, you can change the authentication settings at the virtual directory level.

The authentication settings on virtual directories are different from authentication settings on the Default Web Site and Exchange Back End website. By default, these websites allow anonymous access. This means that anyone can access the server's home page without authenticating themselves. If you disable anonymous access at the server level and enable some other type of authentication, users need to authenticate themselves twice: once for the server and once for the virtual directory they want to access.

The preferred way to manage authentication settings is to use the appropriate cmdlet in Exchange Management Shell:

- For [ActiveSync](#), use [Set-ActiveSyncVirtualDirectory](#)
- For [Autodiscover](#), use [Set-AutodiscoverVirtualDirectory](#)
- For [ECP](#), use [Set-EcpVirtualDirectory](#)
- For [OAB](#), use [Set-OabVirtualDirectory](#)
- For [OWA](#), use [Set-OwaVirtualDirectory](#)
- For [PowerShell](#), use [Set-PowerShellVirtualDirectory](#)
- For [Exchange Web Services](#), use [Set-WebServicesVirtualDirectory](#)

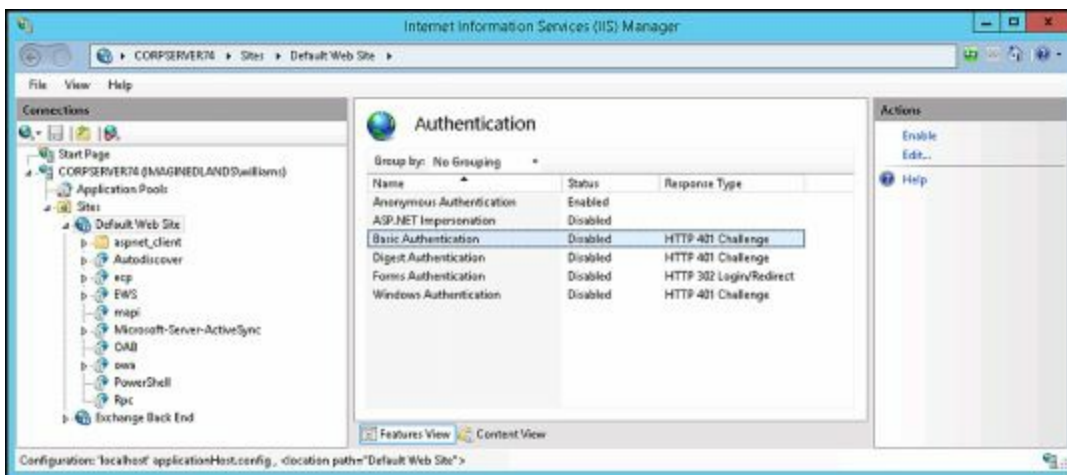
As an example, to disable basic authentication on the default ActiveSync directory, you would enter:

Set-ActiveSyncVirtualDirectory –Identity "TVPRESS\microsoft-server-activesync" –
BasicAuthEnabled \$false

You can change the authentication settings for an entire site or for a particular virtual directory by completing the following steps:

1. Start IIS Manager. In Server Manager, click Tools, and then select Internet Information Services (IIS) Manager.
2. In IIS Manager, navigate to the level you want to manage, and then double-click the Authentication feature. On the Authentication page, you should see the available authentication modes. If a mode you want to use is not available, you need to install and enable the related role service using Server Manager's Add Role Services Wizard.
3. To enable or disable anonymous access, select Anonymous Authentication and then click Enable or Disable as appropriate.

NOTE With anonymous access, IIS uses an anonymous user account for access to the server. The anonymous user account is named IUSR_ServerName, such as IUSR_Mailer1. If you use this account, you don't need to set a password. Instead, let IIS manage the password. If you want to use a different account, click Edit, and then click Set to specify the user name and password for a different account to use for anonymous access.



4. To configure other authentication methods, select the authentication method, and then click Enable or Disable as appropriate. Keep the following in mind:
 - Disabling basic authentication might prevent some clients from accessing resources remotely. Clients can log on only when you enable an authentication method that they support.
 - A default domain isn't set automatically. If you enable Basic authentication, you can choose to set a default domain that should be used when no domain information is supplied during the logon process. Setting the default domain is useful when you want to ensure that clients authenticate properly.
 - With Basic and Digest authentication, you can optionally specify the realm that can be

- accessed. Essentially, a realm is the DNS domain name or web address that will use the credentials that have been authenticated against the default domain. If the default domain and realm are set to the same value, the internal Windows domain name might be exposed to external users during the user name and password challenge/response.
- If you enable ASP.NET Impersonation, you can specify the identity to impersonate. By default, IIS uses pass-through authentication, and the identity of the authenticated user is impersonated. You can also specify a particular user if necessary.
 - If you enable Forms authentication, you can set the logon URL and cookies settings used for authentication.

Throttling Client Access to Servers

Every Mailbox server in your organization is subject to the default throttling policy. Throttling policies are designed to ensure that users aren't intentionally or unintentionally overloading Exchange. Exchange tracks the resources that each user consumes and applies throttling policy to enforce connection bandwidth limits as necessary.

The default policy is set in place when you install your first Mailbox server running a current version of Exchange Server. There is a single default throttling policy for the organization. You can customize the default policy or add additional policies as necessary.

To manage throttling policy, you use Exchange Management Shell and the `Get-ThrottlingPolicy`, `Set-ThrottlingPolicy`, `New-ThrottlingPolicy`, and `Remove-ThrottlingPolicy` cmdlets. Throttling policy applies to:

- Anonymous access
- Exchange Web Services (EWS)
- IMAP
- MAPI
- Microsoft Exchange ActiveSync (EAS)
- Outlook Web App (OWA)
- OWA Voicemail
- POP
- PowerShell
- PowerShell Web Services
- RPC Client Access

With all of these features except PowerShell, you can specify separate settings for the following:

- Maximum concurrency controls the maximum number of connections a user can have at one time, with `$null` removing the limit. The parameters are `AnonymousMaxConcurrency`, `EASMaxConcurrency`, `EWSMaxConcurrency`, `IMAPMaxConcurrency`, `OWAMaxConcurrency`, `POPMaxConcurrency`, and `PowerShellMaxConcurrency` as well as `OWAVoiceMaxConcurrency` for OWA

voicemail, `PsWsMaxConcurrency` for PowerShell Web Services, and `RcaMaxConcurrency` for RPC Client Access.

- Maximum burst controls the amount of time in milliseconds that a user can use an elevated amount of resources before being throttled, with `$null` removing the limit. The parameters are `AnonymousMaxBurst`, `EASMaxBurst`, `EWSMaxBurst`, `IMAPMaxBurst`, `OWAMaxBurst`, `POPMaxBurst`, and `PowerShellMaxBurst` as well as `OWAVoiceMaxBurst` for OWA voicemail, `PsWsMaxBurst` for PowerShell Web Services, and `RcaMaxBurst` for RPC Client Access.

NOTE Each service also has a cutoff balance, such as `AnonymousCutOffBalance`, and a corresponding recharge rate, such as `AnonymousRechargeRate`. Both values are set in milliseconds. Cutoff balance controls the resource consumption limits for a service before a user is completely blocked from performing operations on the related component. Recharge rate controls the rate at which the cutoff balance is recharged. For example, with anonymous access the cut off is 720 seconds (720000 milliseconds) and the recharge rate is 420 seconds (420000 milliseconds). Thus, the maximum amount of time a user can use an anonymous connection is 12 minutes, but after seven minutes of idle time this cutoff value is fully recharged.

With PowerShell you can specify:

- Maximum number of concurrent PowerShell sessions per user using `PowerShellMaxRunspaces`.
- The time period for determining whether the maximum number of run spaces has been exceeded using `PowerShellMaxRunspacesTimePeriod`.
- Maximum number of cmdlets that a user can run in a given interval before their execution is stopped using `PowerShellMaxCmdlets`.
- The time period for determining whether the maximum number of cmdlets has been exceeded using `PowerShellMaxCmdletsTimePeriod`.
- The maximum number of operations allowed to be executed per user with the `PowerShellMaxCmdletQueueDepth`.
- Maximum number of concurrent Remote PowerShell connections for an Exchange tenant organization using `PowerShellMaxTenantConcurrency`.
- Maximum number of concurrent PowerShell sessions that an Exchange tenant organization can have using `PowerShellMaxTenantConcurrency`.

NOTE Maximum concurrency controls the number of user sessions. Maximum cmdlets controls the number of cmdlets in each user session. The two values together are affected by the maximum queue depth allowed. For example, if five user sessions are allowed, and each can run four cmdlets in a given interval, the maximum queue depth to allow this is 20 (5 user session x 4 cmdlets each = 20). Any value less than 20 restricts the number of operations that can be performed in this scenario.

You can get the default throttling policy by entering: `Get-ThrottlingPolicy default*` or `Get-ThrottlingPolicy | where-object {$_.IsDefault -eq $true}`. You can get the throttling

policy applied to a particular user by entering (Get-Mailbox UserAlias).ThrottlingPolicy where UserAlias is the alias for a user, such as:

```
(Get-Mailbox jimj).ThrottlingPolicy | Get-ThrottlingPolicy
```

REAL WORLD You also can use this technique to list the retention policy, address book policy, role assignment policy, or sharing policy associated with a user mailbox (if any). Here are examples:

```
(Get-Mailbox jimj).RetentionPolicy | Get-RetentionPolicy
```

```
(Get-Mailbox jimj).SharingPolicy | Get-SharingPolicy
```

```
(Get-Mailbox jimj).AddressBookPolicy | Get-AddressBookPolicy
```

```
(Get-Mailbox jimj).RoleAssignmentPolicy | Get-RoleAssignmentPolicy
```

You can create a nondefault throttling policy by using the New-ThrottlingPolicy cmdlet. You can then assign the policy to a mailbox by using the -ThrottlingPolicy parameter of the Set-Mailbox and New-Mailbox cmdlets. In the following example, you apply TempUserThrottlingPolicy to AmyG:

```
Set-Mailbox -Identity amyg -ThrottlingPolicy (Get-ThrottlingPolicy  
TempUserThrottlingPolicy)
```

You can modify default and nondefault throttling policies by using Set-ThrottlingPolicy. To have a user go back to the default policy, set the -ThrottlingPolicy parameter to \$null as shown in this example:

```
Set-Mailbox -Identity amyg -ThrottlingPolicy $null
```

You can find all user mailboxes that currently have a particular policy applied by using Get-Mailbox with a where-object filter. In the following example, you look for all user mailboxes that have the TempUserThrottlingPolicy:

```
$p = Get-ThrottlingPolicy TempUserThrottlingPolicy
```

```
Get-Mailbox | where-object {$_.ThrottlingPolicy -eq $p.Identity}
```

To switch multiple users from one policy to another, you can do the following:

```
$op = Get-ThrottlingPolicy TempUserThrottlingPolicy
```

```
$ms = Get-Mailbox | where-object {$_.ThrottlingPolicy -eq $op.Identity}
```

```
$np = Get-ThrottlingPolicy RestrictedUserThrottlingPolicy
```

```
foreach ($m in $ms) {Set-Mailbox $m.Identity -ThrottlingPolicy $np;}
```

You can remove nondefault policies that aren't currently being applied by using Remove-ThrottlingPolicy. Simply enter Remove-ThrottlingPolicy followed by the name of the policy as shown in this example:

```
Remove-ThrottlingPolicy TempUserThrottlingPolicy
```


Optimizing Access for Web and Mobile Clients


When you deploy new Mailbox servers, you'll need to configure internal and external access URLs for client access protocols. While you are working with the related settings, you should also verify that the authentication settings are optimized for your environment. After you optimize these settings, you may need to double-check them as part of routine troubleshooting.

Configuring Access for OAB

Outlook 2010 and later clients can retrieve the offline address book (OAB) from a web distribution point. The default distribution point is the OAB virtual directory on the Default Web Site. Each distribution point has the following three associated properties:

- **PollInterval** The time interval during which the Microsoft Exchange File Distribution service should poll the generation server for new updates (in minutes)
- **ExternalUrl** The URL from which Outlook clients outside the corporate network can access the OAB
- **InternalUrl** The URL from which Outlook clients inside the corporate network can access the OAB

You can configure web distribution points by completing the following steps:

1. In the Exchange Admin Center, select Servers in the Navigation menu, and then select Virtual Directories to view a list of the front-end virtual directories used by Mailbox servers in the Exchange organization.
2. You'll see an entry for each OAB web distribution point. Select the distribution point you want to configure and then select Edit (). This opens the Properties dialog box.
3. Set the desired polling interval using the Polling Interval text box. The default interval is 480 minutes.
4. The current internal and external URLs are listed. If you want to change the current settings, enter the desired internal and external URLs in the text boxes provided. Click Save.

OAB (Default Web Site)

Server:
CORPSERVER74

Last modified time:
11/16/2015 10:45 AM

Polling interval (minutes):
480

Internal URL:
https://corpserver74.imagedlands.com/OAB
This Internal URL refers to the URL from which Outlook clients inside the corporate network can access this virtual directory.

External URL:
https://mail.imagedlands.com/OAB
This External URL refers to the URL from which Outlook clients outside the corporate network can access this virtual directory.

Save Cancel

After you make changes to the OAB directory, you should verify that you can still access the OAB. If you can't access OAB or suspect there is a configuration problem, you can reset the OAB virtual directory by selecting it in the list of virtual directories, selecting Reset, and then confirming the reset by selecting Reset in the warning dialog box. When you reset a virtual directory, Exchange deletes the virtual directory and then recreates it with its default settings. Resetting a directory means any custom settings will be lost.

Configuring Access for OWA

When you install a Mailbox server, the server is configured with a Default Web Site and the virtual directories discussed previously. Through the OWA virtual directory, you can specify different URLs for internal access and external access to OWA. You can also configure various authentication options.

owa (Default Web Site)

general
authentication
features
file access

Use one or more standard authentication methods

- Integrated Windows authentication
- Digest authentication for Windows domain servers
- Basic authentication

Use forms-based authentication

Logon format:


- Domain\user name
- User principal name (UPN)
- User name only

Logon domain:

 Browse..

Save Cancel

You can configure OWA virtual directory URLs and authentication options by completing the following steps:


1. In the Exchange Admin Center, select Servers in the Navigation menu and then select Virtual Directories to view a list of the front-end virtual directories used by Mailbox servers in the Exchange organization.
2. You'll see an entry for each OWA virtual directory available. Select the OWA virtual directory you want to configure, and then select Edit ().
3. In the Properties dialog box, on the General page, the current internal and external URLs are listed. If you want to change the current settings, enter the internal and external URLs you want to use in the text boxes provided.
4. On the Authentication page, forms-based authentication is configured by default with the logon format set to Domain\User Name. Change this configuration only if you have specific requirements that necessitate a change.
5. Click Save to apply your settings.

After you make changes to the OWA directory, you should verify that you can still access Outlook Web App. If you can't access Outlook Web App or suspect there is a configuration problem, you can reset the OWA virtual directory by selecting it in the list of virtual directories, selecting Reset, and then confirming the reset by selecting Reset in the warning dialog box. When you reset a virtual directory, Exchange deletes the virtual directory and then recreates it with its default settings. Resetting a directory means any custom settings will be lost.

Configuring Access for Exchange ActiveSync

When you install a Mailbox server, the server is configured with a Default Web Site that has a virtual directory for Exchange ActiveSync. Through this virtual directory, you can specify different URLs for internal access and external access to Exchange ActiveSync. You also can configure various authentication options.

You can configure the Exchange ActiveSync URLs and authentication options by completing the following steps:

1. In the Exchange Admin Center, select Servers in the Navigation menu and then select Virtual Directories to view a list of the front-end virtual directories used by Mailbox servers in the Exchange organization.
2. You'll see an entry for each virtual directory available. Select the ActiveSync virtual directory you want to configure, and then select Edit ().
3. In the properties dialog box, on the General page, the current internal and external URLs are listed. If you want to change the current settings, enter the internal and external URLs you want to use in the text boxes provided.
4. On the Authentication page, basic authentication is enabled by default and client certificates are ignored. If your organization uses client certificates, you can clear the Basic Authentication check box and then select either Accept Client Certificates or Require Client Certificates as appropriate.

5. Click Save to apply your settings.

Microsoft-Server-ActiveSync (Default Web Site)

general
▶ authentication

SSL enabled:
True

Select the authentication method or methods that this virtual directory accepts. To enable authentication between the Exchange server and a mobile device, either Basic authentication or Client Certificate authentication is required.

Basic authentication
(Requires the use of SSL certificates to encrypt the passwords that are normally sent in clear text)

Client certificate authentication:

Ignore client certificates
 Accept client certificates
 Require client certificates


Save Cancel

After you make changes to the ActiveSync directory, you should verify that you can still access Exchange ActiveSync. If you can't access Exchange ActiveSync or suspect there is a configuration problem, you can reset the ActiveSync virtual directory by selecting it in the list of virtual directories, selecting Reset, and then confirming the reset by selecting Reset in the warning dialog box. When you reset a virtual directory, Exchange deletes the virtual directory and then recreates it with its default settings. Resetting a directory means any custom settings will be lost.

Configuring Access for ECP

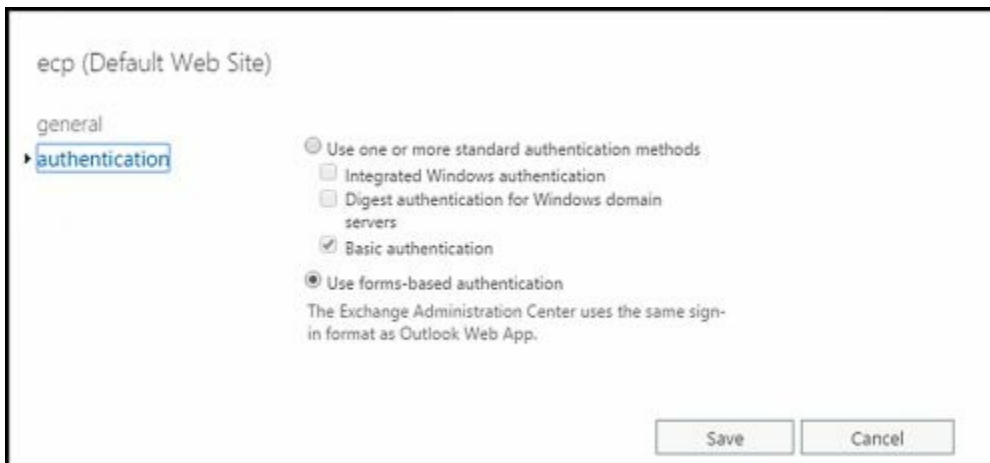
When you install a Mailbox server, the server is configured with a Default Web Site and the virtual directories discussed previously. Through the ECP virtual directory, you can specify different URLs for internal and external access to Exchange Admin Center. You can also configure various authentication options.

You can configure ECP virtual directory URLs and authentication options by completing the following steps:

1. In the Exchange Admin Center, select Servers in the Navigation menu, and then select Virtual Directories to view a list of the front-end virtual directories used by Mailbox servers in the Exchange organization.
2. You'll see an entry for each ECP virtual directory available. Select the ECP virtual directory you want to configure, and then select Edit ().
3. On the General page, the current internal and external URLs are listed. If you want to change the current settings, enter the internal and external URLs you want to use in the text boxes provided.
4. On the Authentication page, basic authentication and forms-based authentication are configured by default. The logon format for forms-based authentication is the

same as the format used for Outlook Web App. Change this configuration only if you have specific requirements that necessitate a change.

5. Click Save to apply your changes.



After you make changes to the ECP directory, you should verify that you can still access the Exchange Admin Center. If you can't access Exchange Admin Center or suspect there is a configuration problem, you can reset the ECP virtual directory by selecting it in the list of virtual directories, selecting Reset, and then confirming the reset by selecting Reset in the warning dialog box. When you reset a virtual directory, Exchange deletes the virtual directory and then recreates it with its default settings. Resetting a directory means any custom settings will be lost.

Chapter 26. Optimizing Client Access Protocols

Mail clients can use a variety of protocols to connect to Exchange server. Generally, Outlook clients will use either MAPI over HTTP or RPC over HTTP while other mail clients will use either POP3 or IMAP4. In this chapter, you'll learn techniques for optimizing these protocols for client access.

Managing RPC and MAPI over HTTP

Outlook clients can use either RPC over HTTP or MAPI over HTTP for connecting to their Exchange mailboxes, with MAPI over HTTP as the preferred option. Generally, these features are enabled and configured automatically when you install Exchange services and no additional configuration is required (see Chapter 13 for exceptions). RPC and MAPI over HTTP also are secure by default, so unauthenticated requests from Outlook clients are blocked from accessing Exchange Server.

Working with RPC and MAPI over HTTP

The only requirement for RPC and MAPI over HTTP is that Exchange servers have properly configured SSL certificates. Because RPC and MAPI over HTTP requests use HTTPS, you must allow port 443 through your firewall. If you already use Outlook Web App with SSL or Exchange ActiveSync with SSL, port 443 should already be open and you do not have to open any additional ports.

As with other services, RPC and MAPI over HTTP have front-end components and back-end components on Mailbox servers:

- [RPC over HTTP uses the Rpc virtual directory on the Default Web Site and the Rpc, RpcProxy and RpcWithCert virtual directories on the Exchange Back End website.](#)
- [MAPI over HTTP uses the Mapi virtual directory on the Default Web Site and the Mapi virtual directory on the Exchange Back End website.](#)

You can use the `Get-OutlookAnywhere` cmdlet to list configuration details for RPC over HTTP. If you use the `-Server` parameter, you can limit the results to a specific server. If you use the `-Identity` parameter, you can examine a particular virtual directory on a server. Listing 26-1 provides the syntax, usage and sample output.

LISTING 26-1 `Get-OutlookAnywhere` cmdlet syntax and usage

Syntax

```
Get-OutlookAnywhere [-Server ServerName] [-DomainController DCName]
```

```
Get-OutlookAnywhere [-Identity VirtualDirId] [-DomainController DCName]
```

Usage

```
Get-OutlookAnywhere
```

```
Get-OutlookAnywhere -Server "MailServer42"
```

```
Get-OutlookAnywhere -Identity "MailServer42\Rpc (Default Web Site)"
```

Sample Output

```
RunspaceId           : 035f41f8-7f92-4b3d-ac89-822b885de085
ServerName           : MailServer42
```

```

SSLOffloading           : True
ExternalHostname       :
InternalHostname       : mailserver42.imaginedlands.local
ExternalClientAuthenticationMethod : Negotiate
InternalClientAuthenticationMethod : Ntlm
IISAuthenticationMethods : {Basic, Ntlm, Negotiate}
XropUrl                 :
ExternalClientsRequireSsl : False
InternalClientsRequireSsl : False
MetabasePath           : IIS://CorpServer74.imaginedlands.local/W3SVC/1/ROOT/Rpc
Path : C:\Program Files\Microsoft\Exchange Server\V15\FrontEnd\HttpProxy\rpc
ExtendedProtectionTokenChecking : None
ExtendedProtectionFlags : {}
ExtendedProtectionSPNList : {}
AdminDisplayVersion    : Version 15.1 (Build 225.42)
Server                 : CORPSEVER74
AdminDisplayName       :
ExchangeVersion        : 0.20 (15.0.0.0)
Name                   : Rpc (Default Web Site)
...
Id                     : CORPSEVER74\Rpc (Default Web Site)
OriginatingServer     : CorpServer91.imaginedlands.local
IsValid                : True
ObjectState            : Changed

```

Listing 26-2 provides the syntax, usage, and sample output for Get-MapiVirtualDirectory cmdlet, which you can use to list configuration details for MAPI over HTTP. When using this cmdlet, you can use the `-Server` parameter to limit the results to a specific server or the `-Identity` parameter to examine a particular virtual directory on a server. To get a complete listing, be sure to format the output as shown in the examples.

LISTING 26-2 Get-MapiVirtualDirectory cmdlet syntax and usage

Syntax

```
Get-MapiVirtualDirectory [-Server ServerName] [-DomainController DCName]
```

```
Get-MapiVirtualDirectory [-Identity VirtualDirId] [-DomainController DCName]
```

Usage

```
Get-MapiVirtualDirectory |fl
```

```
Get-MapiVirtualDirectory -Server "MailServer42" |fl
```

```
Get-MapiVirtualDirectory -Identity "MailServer42\Rpc (Default Web Site)" | fl
```

Sample Output


```
RunspaceId           : 035f41f8-7f92-4b3d-ac89-822b885de085
```

IISAuthenticationMethods : {Ntlm, OAuth, Negotiate}
MetabasePath : IIS://CorpServer74.imaginedlands.local/W3SVC/1/ROOT/mapi
Path : C:\Program Files\Microsoft\Exchange Server\V15\FrontEnd\HttpProxy\mapi
ExtendedProtectionTokenChecking : None
ExtendedProtectionFlags : {}
ExtendedProtectionSPNList : {}
AdminDisplayVersion : Version 15.1 (Build 225.42)
Server : CORPSEVER74
InternalUrl : https://corpserver74.imaginedlands.local/mapi
InternalAuthenticationMethods : {Ntlm, OAuth, Negotiate}
ExternalUrl :
ExternalAuthenticationMethods : {Ntlm, OAuth, Negotiate}
AdminDisplayName :
ExchangeVersion : 0.10 (14.0.100.0)
Name : mapi (Default Web Site)
...
Id : CORPSEVER74\mapi (Default Web Site)
OriginatingServer : CorpServer91.imaginedlands.local
IsValid : True
ObjectState : Changed

Configuring URLs and Authentication

When you install a Mailbox server, the server is configured with a Default Web Site and the virtual directories discussed previously. Through the Rpc and Mapi virtual directories on the front-end, you can specify different URLs for internal and external access to RPC and MAPI over HTTP. You can also configure various authentication options.

Although a graphical interface for setting Mapi virtual direction options isn't available, you can configure RPC virtual directory URLs and authentication options in Exchange Admin Center. To do this, complete the following steps:

1. Select Servers in the Navigation menu, and then select the Servers tab to view a list of servers in the Exchange organization.
2. Select the server with which you want to work, and then select Edit ().
3. When you select the Outlook Anywhere page in the Properties dialog box, the current internal and external URLs are listed. If you want to change the current settings, enter the internal and external URLs you want to use in the text boxes provided.



4. Select an available external authentication method. You can select Basic Authentication, NTLM Authentication, or Negotiate. Although NT LAN Manager (NTLM) authentication is more secure than basic authentication, the most secure option is Negotiate, which configures Outlook Anywhere to use Integrated Windows Authentication.
5. Select the Allow Secure Channel (SSL) Offloading check box only if you have configured an advanced firewall server to work with Exchange 2016 and handle your SSL processing.
6. Click Save to apply your settings.

You also can use the Set-OutlookAnywhere cmdlet to modify the RPC over HTTP configuration. See Listing 26-3 for syntax and usage. The `-IISAuthenticationMethods` parameter sets the authentication method for the `/rpc` virtual directory as either Basic, NTLM or Negotiate and disables all other methods. The `-ExternalClientAuthenticationMethod` and `-InternalClientAuthenticationMethod` parameters set permitted authentication methods for external and internal clients respectively. You can also control whether SSL is required for external and internal clients using the `-ExternalClientRequireSsl` and `-InternalClientRequireSsl` parameters respectively.

LISTING 26-3 Set-OutlookAnywhere cmdlet syntax and usage

Syntax

```
Set-OutlookAnywhere -Identity VirtualDirId [-DomainController DCName]
[-DefaultAuthenticationMethod {AuthMethod}]
[-ExternalClientAuthenticationMethod {AuthMethod}]
[-ExternalClientRequireSsl {$true|$false}]
[-ExternalHostName ExternalHostName]
[-IISAuthenticationMethods <Basic | NTLM | Negotiate>]
[-InternalClientAuthenticationMethod {AuthMethod}]
[-InternalClientRequireSsl {$true|$false}]
[-InternalHostName InternalHostName]
[-Name Name]
[-SSLOffloading <$true | $false>]
```



```
{AuthMethod}  
<Basic | Digest | NTLM | Fba | WindowsIntegrated | LiveIdFba |  
LiveIdBasic | WSSecurity | Certificate | NegoEx | OAuth | Adfs | Kerberos  
| Negotiate | LiveIdNegotiate | Misconfigured>
```

Usage

```
Set-OutlookAnywhere -Identity "CorpSvr127\Rpc (Default Web Site)"  
-ExternalHostName "mail.tvpress.com"  
-InternalHostName "mailserver21.tvpress.com"  
-ExternalAuthenticationMethod "Negotiate"  
-SSLOffloading $true
```

To configure MAPI virtual directory URLs and authentication options, you'll use the Set-MapiVirtualDirectory cmdlet. Listing 26-4 provides the syntax and usage. Use the -ExternalUrl and -InternalUrl parameters to set the external and internal Mapi URLs respectively.

LISTING 26-4 Set-MapiVirtualDirectory cmdlet syntax and usage

Syntax

```
Set-MapiVirtualDirectory -Identity VirtualDirId [-DomainController DCName]  
[-IISAuthenticationMethods <Basic | NTLM | Negotiate>]  
[-InternalUrl Url]  
[-ExternalUrl Url]
```

Usage

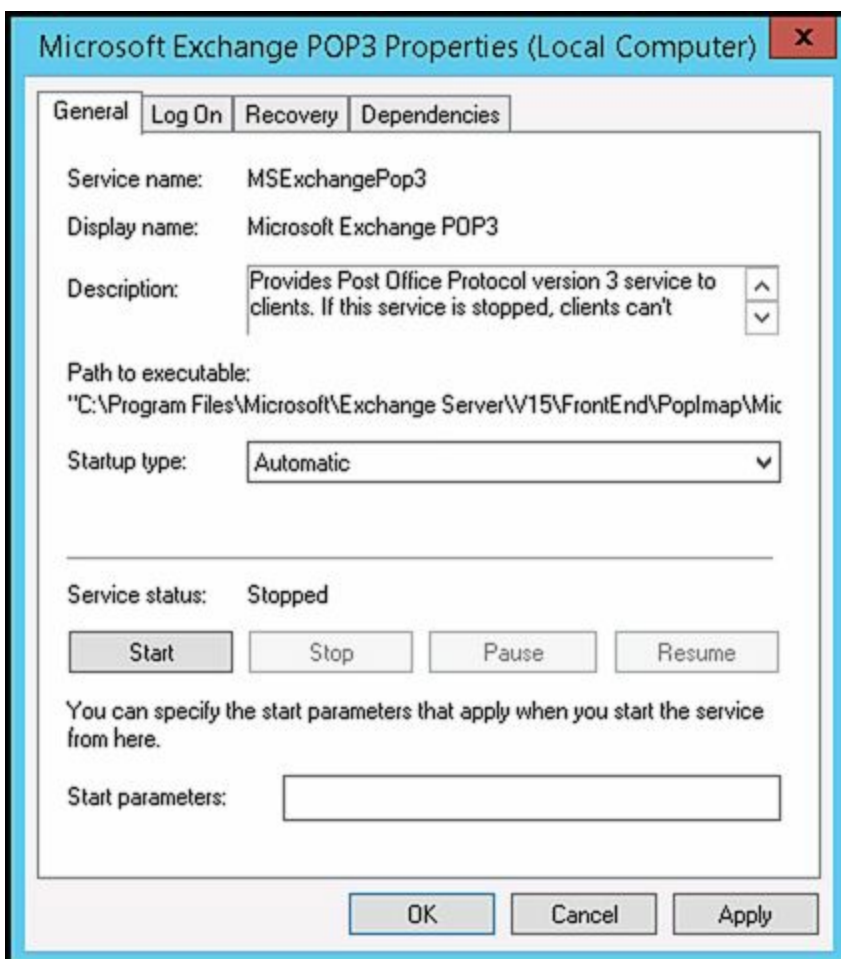
```
Set-MapiVirtualDirectory -Identity "CorpSvr127\Mapi (Default Web Site)"  
-InternalUrl "http://mailserver21.imaginedlands.com/mapi"  
-ExternalUrl "http://mail.imaginedlands.com/mapi"
```

Enabling the POP3 and IMAP4 Services

Clients that retrieve mail using POP3 or IMAP4 send mail using SMTP. SMTP is the default mail transport in Exchange Server 2016. To enable POP3 and IMAP4, you must first start the POP3 and IMAP4 services on the Exchange servers that will provide these services. You must then configure these services to start automatically in the future. You should also review the related settings for each service and make changes as necessary to optimize the way these services are used in your Exchange organization.

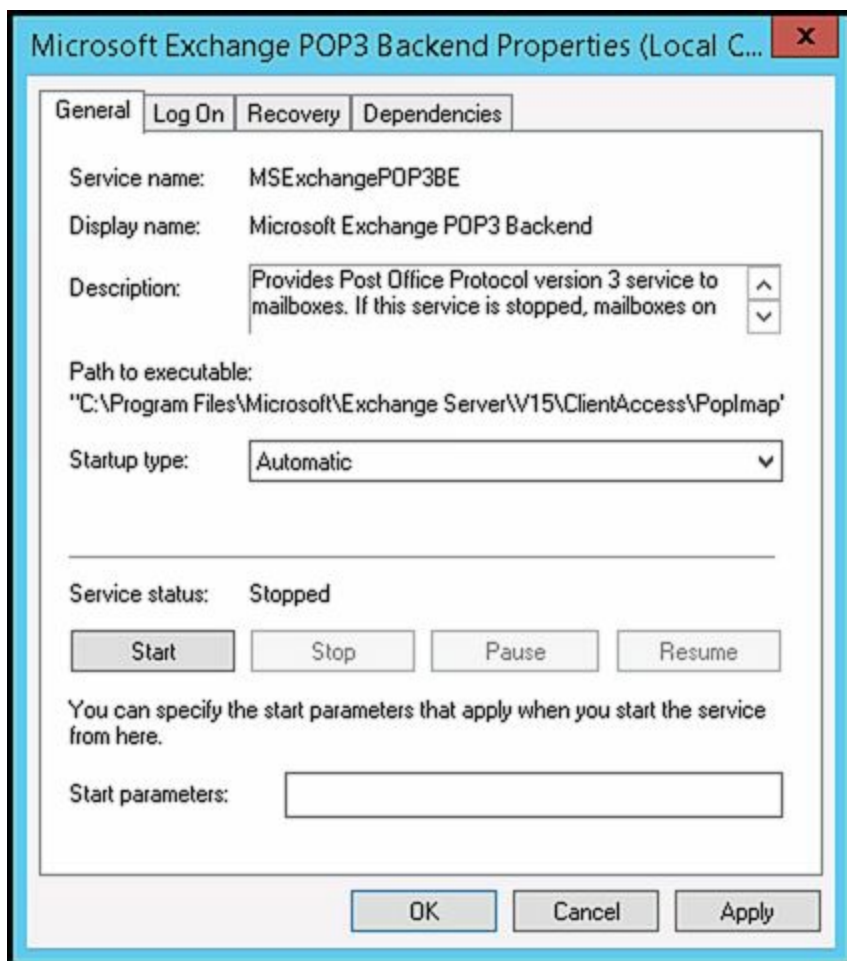
Because the client access infrastructure has two-layers with a front-end component and a back-end component, there are corresponding front-end and back-end services for both POP3 and IMAP4 running on Mailbox servers. You can enable and configure the front-end POP3 service for automatic startup by completing these steps:

1. Start the Services utility. In Server Manager, click Tools, and then select Services.
2. Right-click Microsoft Exchange POP3, and then select Properties.
3. On the General tab, under Startup Type, select Automatic and then click Apply.
4. Under Service Status, click Start, and then click OK.

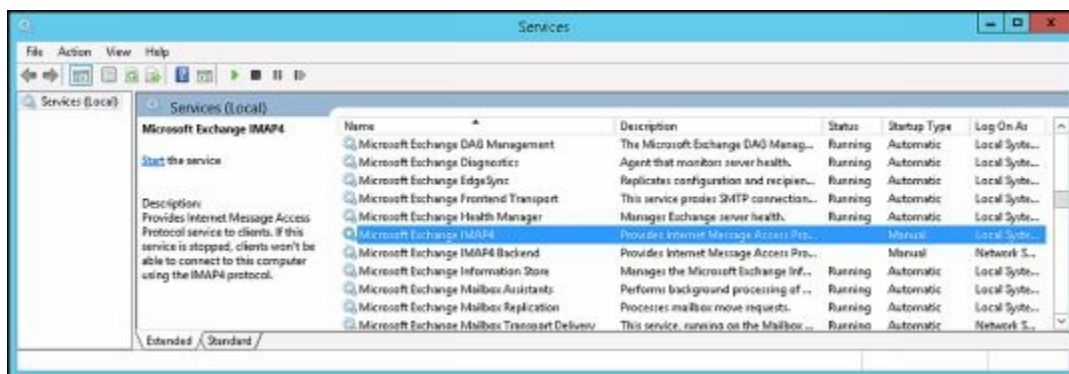


The corresponding back-end service is the POP3 Backend service. You can enable and configure the POP3 Backend service for automatic startup by completing the following steps:

1. Start the Services utility. In Server Manager, click Tools, and then select Services.
2. Right-click Microsoft Exchange POP3 Backend, and then select Properties.
3. On the General tab, under Startup Type, select Automatic and then click Apply.
4. Under Service Status, click Start, and then click OK.



If you want to enable IMAP4, configure the Microsoft Exchange IMAP4 service as well as the Microsoft Exchange IMAP4 Backend service on your Mailbox servers for automatic startup and then start the services. Use the same techniques as discussed previously for POP3.



You can use the Set-Service cmdlet to enable and configure POP3 and IMAP4 as well. The identifiers for the services are as follows:

- MExchangePop3 for the POP3 front-end service

- [MSExchangePop3BE](#) for the POP3 back-end service
- [MSExchangeIMAP4](#) for the IMAP4 front-end service
- [MSExchangeIMAP4BE](#) for the IMAP4 back-end service

Use the `–StartupType` parameter to set the startup type as Automatic, Manual, or Disabled. Use the `–Status` parameter to set the status as Running, Paused, or Stopped. The following examples enable POP3 and IMAP4 for automatic startup and then start the services:

```
Set-Service –Name MSExchangePop3 –StartupType Automatic –Status Running
```

```
Set-Service –Name MSExchangeImap4 –StartupType Automatic –Status Running
```

The following examples enable the POP3 and IMAP4 back end services for automatic startup, and then start the services:

```
Set-Service –Name MSExchangePop3BE –StartupType Automatic –Status Running
```

```
Set-Service –Name MSExchangeImap4BE –StartupType Automatic –Status Running
```

Optimizing POP3 and IMAP4 Settings

As alternatives to RPC and MAPI over HTTP, Exchange 2016 supports Internet Message Access Protocol 4 (IMAP4) and Post Office Protocol 3 (POP3). IMAP4 is a protocol for reading mail and accessing public and private folders on remote servers. Clients can log on to an Exchange server and use IMAP4 to download message headers and then read messages individually while online. POP3 is a protocol for retrieving mail on remote servers. Clients can log on to an Exchange server and then use POP3 to download their mail for offline use.

By default, POP3 (version 3) and IMAP4 (rev 1) are configured for manual startup. Because other client access protocols offer so much more than POP and IMAP, they are the preferred way for clients to access Exchange Server. That said, if you still have users who want to use POP3 and IMAP4 to access Exchange Server, you can configure this, but you should try to move these users to other client access protocols.

As you configure POP3 and IMAP4 access don't forget that the client access infrastructure has two layers:

- A front end that you can customize to control the way users access and work with POP3 and IMAP4
- A back end that handles the back-end processing but that you only modify to control the options that the front end uses for working with the back-end processes

Thus, although you typically modify the front-end settings for POP3 and IMAP4 to customize the environment for users, you rarely modify the related back-end components.

Configuring POP3 and IMAP4 Bindings

POP3 and IMAP4 have related IP address and TCP port configuration settings. The default IP address setting is to use any available IP address. On a multihomed server, however, you'll usually want messaging protocols to respond on a specific IP address in which case you need to change the default setting.


The default port setting depends on the messaging protocol being used and whether SSL is enabled or disabled. For users to be able to retrieve mail using POP3 and IMAP4, you must open the related messaging ports on your organization's firewalls. The default port settings for key protocols used by Exchange Server 2016 are as follows:

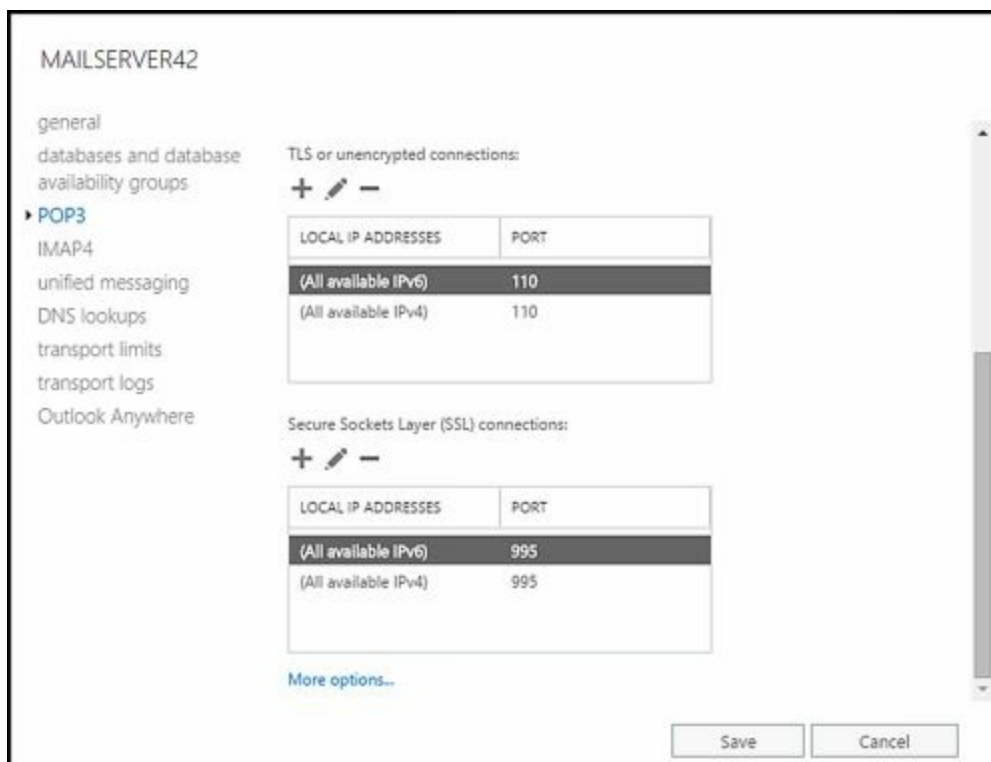
- With SMTP, the default port is 25 and the default secure port is 587.
- With HTTP, the default port is 80 and the default secure port is 443.
- With POP3, the default port is 110 and the default secure port is 995.
- With IMAP4, the default port is 143 and the default secure port is 993.

In Exchange Management Shell, you can manage POP3 and IMAP4 by using the following cmdlets:

- **Get-POPSettings** Lists POP3 configuration settings
- **Set-POPSettings** Configures POP3 settings
- **Test-POPConnectivity** Tests the POP3 configuration
- **Get-IMAPSettings** Lists IMAP4 configuration settings
- **Set-IMAPSettings** Configures IMAP4 settings
- **Test-IMAPConnectivity** Tests the IMAP4 configuration

The bindings for POP3 and IMAP4 use a unique combination of an IP address and a TCP port. To change the IP address or port number for POP3 or IMAP4, complete the following steps:



1. In the Exchange Admin Center, select Servers in the Navigation menu, and then select the Servers tab to view a list of servers in the Exchange organization.
2. Select the server with which you want to work, and then select Edit ().
3. In the Properties dialog box, select the POP3 or IMAP4 page as appropriate for the service you want to configure.
4. If you scroll down, you'll see the currently assigned IP addresses and ports used for TLS or unencrypted connections and SSL connections. The default configuration is as follows: POP3 and IMAP4 are configured to use all available IPv4 and IPv6 addresses, POP3 uses port 110 for TLS or unencrypted connections and port 995 for SSL connections, and IMAP4 uses port 143 for TLS or unencrypted connections and port 993 for SSL connections.



5. To configure IP addresses and ports for TLS or unencrypted connections, use the following options on the TLS Or Unencrypted Connections panel:




- Add () Adds a TCP port on a per-IP address basis or all unassigned IP address

basis. Click Add, and then specify the IP address and port you want to use.

- Edit () Allows you to edit the IP address and port settings for the currently selected entry in the Address list box.
- Remove () Allows you to remove the IP address and port settings for the currently selected entry in the Address list box.

NOTE The IP address/TCP port combination must be unique. You can assign the same port as long as the protocol is configured to use a different IP address. You can also assign the same IP address and use a different port.

6. To configure IP addresses and ports for secure connections, use the following options on the Secure Sockets Layer (SSL) Connections panel:

- Add () Adds a TCP port on a per-IP address basis or an all-unassigned IP address basis. Click Add, and then specify the IP address and port you want to use.
- Edit () Allows you to edit the IP address and port settings for the currently selected entry in the Address list box.
- Remove () Allows you to remove the IP address and port settings for the currently selected entry in the Address list box.


7. Click Save to apply your settings. When you add new ports, you must open the related messaging ports on your organization's firewalls.

8. Use the Services utility to restart the Exchange POP3 or IMAP4 service. Restarting the service applies the new settings.

Configuring POP3 and IMAP4 Authentication

By default, POP3 and IMAP4 clients pass connection information and message data through a secure TLS connection. A secure TLS connection requires the Exchange servers to have properly configured SSL certificates with POP3, IMAP4, or both as assigned services.

Secure TLS connections are the best option to use when corporate security is a high priority and secure communication channels are required. That said, you have two other options for configuring communications: plain-text authentication and logon using integrated Windows authentication. Configure these authentication options by completing the following steps:

1. In the Exchange Admin Center, select Servers in the Navigation menu, and then select the Servers tab to view a list of servers in the Exchange organization.
2. Select the server with which you want to work, and then select Edit ().
3. In the Properties dialog box, select the POP3 or IMAP4 page as appropriate for the service you want to configure.

4. For Logon Method, do one of the following, and then click Save:
 - Select **Basic Authentication (Plain text)** to use unsecure plain text for communications.
 - Select **Integrated Windows Authentication (Plain text)** to use secure communications with Windows authentication.
5. Use the Services utility to restart the Exchange POP3 or IMAP4 service. Restarting the service applies the new settings.

When Exchange is configured to use SSL and secure TLS connections, mail clients also should be configured to use either SSL or secure TLS. You can configure an Outlook client to use SSL or secure TLS by completing the following steps:

1. Do one of the following:
 - In Office 2010, click the **Office** button, click **Account Settings**, and then select the **Account Settings** option.
 - In Office 2013 or Office 2016, click the **File** tab. Next, select the **Account Settings** option and then select **Account Settings**.
2. In the Account Settings dialog box, select the POP3/IMAP4 account, and then click **Change**.
3. In the Change E-Mail Account dialog box, click **More Settings**.
4. On the **Advanced** tab in the Internet E-Mail Settings dialog box, select **SSL, TLS** or **Auto** as the type of encrypted connection.
5. Click **OK**. Click **Next**, and then click **Finish**. Click **Close**.

Configuring Connection Settings for POP3 and IMAP4


You can control incoming connections to POP3 and IMAP4 in two ways. You can set a limit on the number of simultaneous connections, and you can set a connection time-out value.

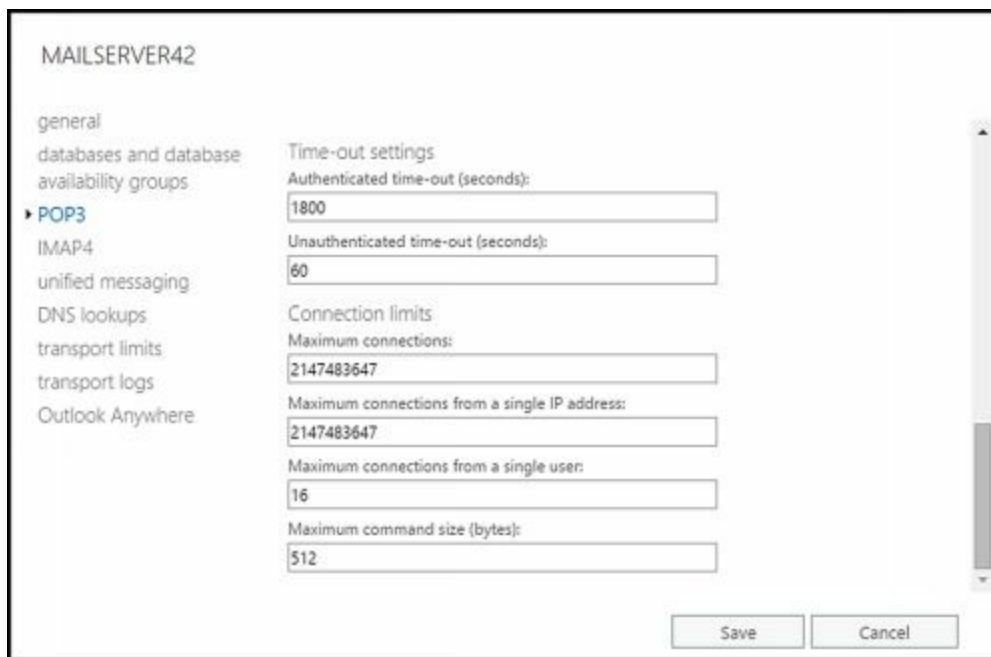
POP3 and IMAP4 normally accept a maximum of 2,147,483,467 connections each and a maximum of 16 connections from a single user, and in most environments these are acceptable settings. However, when you're trying to prevent the underlying server hardware from becoming overloaded or you want to ensure resources are available for other features, you might want to restrict the number of simultaneous connections to a much smaller value. When the limit is reached, no other clients are permitted to access the server. The clients must wait until the connection load on the server decreases.

The connection time-out value determines when idle connections are disconnected. Normally, unauthenticated connections time out after they've been idle for 60 seconds and authenticated connections time out after they've been idle for 1,800 seconds (30 minutes). In most situations, these time-out values are sufficient. Still, at times you'll want to increase the time-out values, and this primarily relates to clients who get disconnected when downloading large files. If you discover that clients are being disconnected during large downloads, the time-out values are one area to examine. You'll also want to look at the maximum command size. By default, the maximum

command size is restricted to 512 bytes.

You can modify connection limits and time-outs by completing the following steps:

1. In the Exchange Admin Center, select Servers in the Navigation menu, and then select the Servers tab to view a list of servers in the Exchange organization.
2. Select the server with which you want to work, and then select Edit ().
3. In the Properties dialog box, select the POP3 or IMAP4 page as appropriate for the service you want to configure.
4. Scroll down and then click More Options to display the additional configuration options.



MAILSERVER42

general
databases and database
availability groups
► POP3
IMAP4
unified messaging
DNS lookups
transport limits
transport logs
Outlook Anywhere

Time-out settings
Authenticated time-out (seconds):
1800
Unauthenticated time-out (seconds):
60

Connection limits
Maximum connections:
2147483647
Maximum connections from a single IP address:
2147483647
Maximum connections from a single user:
16
Maximum command size (bytes):
512

Save Cancel


5. To set time-out values for authenticated and unauthenticated connections, enter the desired values in the Authenticated Time-Out and Unauthenticated Time-Out text boxes, respectively. The valid range for authenticated connections is from 30 through 86,400 seconds. The valid range for unauthenticated connections is from 10 through 3,600 seconds.
6. To set connection limits, enter the desired limits in the text boxes on the Connection Limits panel. The valid input range for maximum connections is from 1 through 2,147,483,467. The valid input range for maximum connections from a single IP address is from 1 through 2,147,483,467. The valid input range for maximum connections from a single user is from 1 through 2,147,483,467. The valid input range for maximum command size is from 40 through 1,024 bytes.
7. Click Save to apply your settings. Use the Services utility to restart the Exchange POP3 or IMAP4 service. Restarting the service applies the new settings.

Configuring Message Retrieval Settings for POP3 and IMAP4

Message retrieval settings for POP3 and IMAP4 control the following options:

- **Message formatting** Message format options allow you to set rules that POP3 and IMAP4 use to format messages before clients read them. By default, when POP3 or IMAP4 clients retrieve messages, the message body is converted to the best format for the client and message attachments are identified with a Multipurpose Internet Mail Extensions (MIME) content type based on the attachment's file extension. You can change this behavior by applying new message MIME formatting rules. Message MIME formatting rules determine the formatting for elements in the body of a message. Message bodies can be formatted as plain text, HTML, HTML and alternative text, enriched text, or enriched text and alternative text.
- **Message sort order** Message sort order options allow you to control the time sorting of messages during new message retrieval. By default, POP3 sorts messages in ascending order according to the time/date stamp. This ensures that the most recent messages are listed first. You can also sort messages by descending order, which places newer messages lower in the message list.

You can modify message retrieval settings by completing the following steps:

1. In the Exchange Admin Center, select Servers in the Navigation menu, and then select the Servers tab to view a list of servers in the Exchange organization.
2. Select the server with which you want to work, and then select Edit ().
3. In the Properties dialog box, select the POP3 or IMAP4 page as appropriate for the service you want to configure.

The screenshot shows the 'Properties' dialog box for server 'CORPSEVER74'. The left-hand navigation pane lists various categories, with 'POP3' selected. The main area displays configuration options for the POP3 service:

- Message MIME format:** A dropdown menu currently set to 'Best body format'.
- Message sort order:** A dropdown menu currently set to 'Ascending'.
- Logon method:** A dropdown menu currently set to 'Secure TLS connection'.
- Banner string:** A text input field containing the text 'The Microsoft Exchange POP3 service is ready.'

At the bottom of the dialog, there are 'Save' and 'Cancel' buttons.

4. Use the Message MIME Format list to choose the desired body format for messages. As discussed previously, the options are Text, HTML, HTML And Alternative Text, Enriched Text, Enriched Text And Alternative Text, or Best Body Format.
5. If you are working with POP3, use the Message Sort Order list to specify the default sort order for message retrieval. Select Descending for descending sort order during message retrieval or Ascending for ascending sort order.
6. Click Save to apply your settings. Use the Services utility to restart the Exchange POP3 or IMAP4 service. Restarting the service applies the new settings.

Chapter 27. Configuring Mobile Messaging

In our increasingly connected world, users want to be able to access email, calendars, contacts, and scheduled tasks no matter the time or place. With Microsoft Exchange 2016 and Microsoft Exchange Online, you can make anywhere, anytime access to Exchange data a real possibility. Exchange 2016 and Exchange Online support wireless access for users with many types of mobile devices via Exchange ActiveSync and Outlook Web App for Devices.

Exchange ActiveSync allows users to link mobile devices to their Exchange accounts so that Exchange synchronizes mail data with the mobile device. Because mail and other data is stored on the device users can access their email, calendar, contacts, and scheduled tasks whether they are online or offline.

Outlook Web App for Devices allows users to access Outlook Web App on a tablet or smartphone simply by accessing the app in the device's browser and logging in. Unlike Exchange ActiveSync, Outlook Web App for Devices does not normally store mail and related data in a file cache on a user's mobile device.

Mobile access to Exchange Server is supported on smart phones and other mobile devices. Most mobile devices include extensions that permit the use of additional features, including

- [Autodiscover](#)
- [Direct Push](#)
- [Remote Device Wipe](#)
- [Password Recovery](#)
- [Direct File Access](#)
- [Remote File Access](#)

In Exchange Server, these features are all enabled by default. The sections that follow discuss how these features work and how related options are configured. Many additional options are available for fine-tuning the mobile access configuration as well, including mobile device mailbox policy and Outlook Web App policy. These policies are discussed as well.

Mastering Mobile and Wireless Access Essentials

Because sensitive data might be stored on a user's mobile device with Exchange ActiveSync, several safeguards are in place to prevent unauthorized access to this data. The first safeguard is a device password, which can be reset remotely by the user or by an administrator. The second safeguard is a remote wipe feature that remotely instructs a mobile device to delete all its Exchange and corporate data. A third safeguard is a data encryption requirement, which can be enabled and enforced.

Getting Started with Exchange ActiveSync

When you install Exchange 2016 or use Exchange Online, Exchange ActiveSync and Outlook Web App for Devices are automatically configured for use, which makes these features easy to manage. However, there are still some essential concepts you should know to manage them more effectively. This section explains these concepts.

Exchange ActiveSync allows users with smart phones and other mobile devices to initiate synchronization with Exchange to keep their data up to date and receive notices from Exchange that trigger synchronization through the Direct Push feature. *Direct Push* is a key feature about which you probably want to know a bit more. Direct Push works like this:

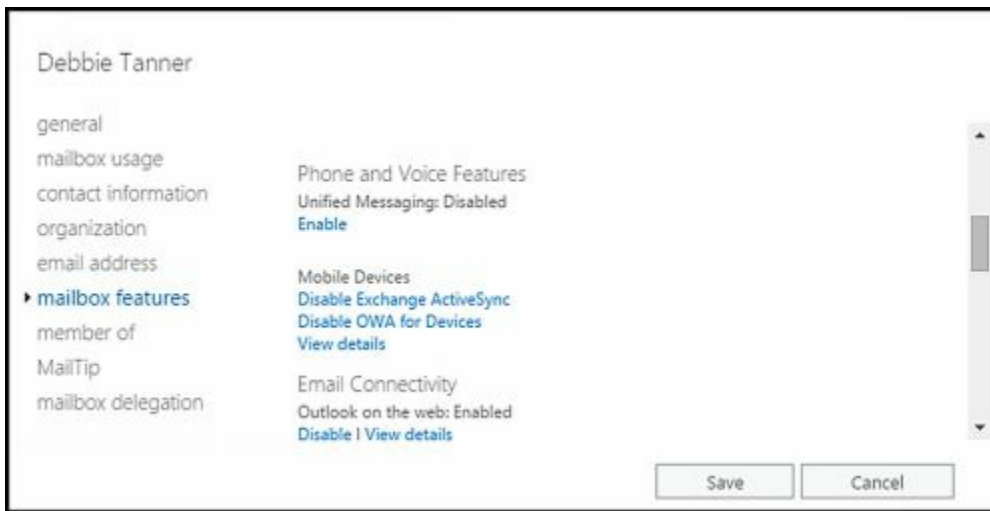
1. The user configures her mobile device to synchronize with Exchange, selecting specific Exchange folders that she wants to keep up to date.
2. When a new message arrives in a designated sync folder, a control message is sent to the mobile device.
3. The control message initiates a data synchronization session, and the device performs background synchronization with Exchange.

After synchronization, users can then access their data while they are offline. In Exchange 2016, Direct Push is either enabled or disabled as is Exchange ActiveSync itself. Because Direct Push uses HTTPS, TCP port 443 must be open on your firewall between the Internet and the Mailbox server to which the user is connecting.

Managing ActiveSync and OWA for Devices

With Exchange Online, Exchange ActiveSync and Outlook Web App for Devices are enabled by default and you cannot change this setting. With Exchange 2016, Exchange ActiveSync is enabled for each user by default, but you can disable Exchange ActiveSync for specific users as necessary by completing the following steps:

1. In Exchange Admin Center, select Recipients in the Navigation menu, and then select Mailboxes.
2. You should now see a list of users with Exchange mailboxes in the organization. Double-click the user's name to open the Properties dialog box for the user account.



3. On the Mailbox Features page, the enabled mobile and web access features for the user are displayed.

- To disable Exchange ActiveSync for this user, under Mobile Devices, select Disable Exchange ActiveSync, and then click Yes.
- To enable Exchange ActiveSync for this user, under Mobile Devices, select Enable Exchange ActiveSync, and then click Yes.
- To disable OWA For Devices for this user, under Mobile Devices, select Disable OWA For Devices, and then click Yes.
- To enable OWA For Devices for this user, under Mobile Devices, select Enable OWA For Devices, and then click Yes.

4. Click Save.

REAL WORLD Exchange ActiveSync notifications are sent over the Internet. The actual process of receiving synchronization requests and sending synchronization notifications is handled by Exchange. Exchange ActiveSync is, in fact, configured as an ASP.NET application on the web server. For Exchange ActiveSync to work properly, IIS server must be configured properly.

To define organization-wide security and authentication options, you can use mobile device mailbox policies. When you install Exchange 2016 or use Exchange Online, a default mobile device mailbox policy is created. Through mobile device mailbox policy settings, you can precisely control mobile browsing capabilities for all users in the enterprise, including:

- Whether Apple mobile devices can get push notifications
- Whether passwords are required, and how passwords must be configured
- Synchronization settings to include in addition to calendar and email items
- Permitted devices and device options, such as whether a device can use Wi-Fi, infrared, Bluetooth, storage cards, or its built-in camera
- Whether the device, its storage cards, or both must be encrypted

Although you configure many mobile device settings in Exchange Admin Center, you will need to use Exchange Admin Shell to fully customize mobile device options.

Configuring Autodiscover

The Autodiscover service simplifies the provisioning process for mobile devices and for Outlook 2010 and later clients by returning the required Exchange settings after a user enters his or her email address and password. This provisioning eliminates the need to configure mobile carriers in Exchange Server, as well as the need to download and install the carriers list on mobile devices.

Understanding Autodiscover

Autodiscover is enabled by default. The Default Web Site associated with a particular Mailbox server has an associated Autodiscover virtual directory that handles proxying and authentication for Autodiscover. The Exchange Back End website associated with the Mailbox server hosting the user's mailbox has an Associated Autodiscover virtual directory through which devices can be provisioned. These virtual directories handle Autodiscover requests:

- Whenever an Outlook client queries for service details
- Whenever a user account is configured or updated
- Whenever the network connection changes

Each Mailbox server is configured with a service connection point that contains an authoritative list of Autodiscover URLs for the associated Active Directory forest. The Autodiscover service URL for the service connection point is either `https://SMTPdomain/autodiscover/autodiscover.xml` or `https://autodiscover.SMTPdomain/autodiscover/autodiscover.xml`, where *SMTPdomain* is the name of the SMTP domain to which the client wants to connect, such as Imaginedlands.com.

ExServerName is the name of a Mailbox server in the site to which the client is connecting. For example, if the user's email address is `tony@contoso.com`, the primary SMTP domain address is `contoso.com`.

When the client connects to Active Directory, the client authenticates to Active Directory by using the user's credentials and then queries for the available service connection point objects. One service connection point object is created for each Mailbox server deployed in the Exchange organization. This object contains a `ServiceBindingInfo` attribute with the fully qualified domain name of the corresponding Mailboxserver in the form `https://ServerFQDN/autodiscover/autodiscover.xml`, where *ServerFQDN* is the fully qualified name of the Mailbox server. After the client obtains and enumerates the service connection point instances, the client connects to the first Mailbox server in the enumerated list and obtains the profile information needed to connect to the user's mailbox. This profile is formatted with XML and also includes a list of available Exchange features.

Maintaining Autodiscover

When you install a Mailbox server, the server is configured with a Default Web Site that has a virtual directory for Autodiscover. Through this virtual directory, you can specify different URLs for internal access and external access to Autodiscover. You also can configure various authentication options.

In the Exchange Admin Center, select Servers in the Navigation menu and then select Virtual Directories to view a list of the front-end virtual directories used by Mailbox servers in the Exchange organization, which includes an entry for each Autodiscover virtual directory available. If you've made any changes to an Autodiscover virtual directory, you should verify that you can still access Autodiscover. If you can't access Autodiscover or suspect there is a configuration problem, you can reset the Autodiscover virtual directory by selecting it in the list of virtual directories, and then selecting Reset. In the Warning dialog box, enter the full file path to a network share in which a settings file can be created to store the current settings for the Autodiscover virtual directory, such as `\\mailserver21\updates\Autodiscoverlog.txt`. Finally, confirm that you want to reset the virtual directory by selecting Reset. When you reset a virtual directory, Exchange deletes the virtual directory and then recreates it with its default settings. Resetting a directory means any custom settings will be lost. To complete the process, you must run the `iisreset /noforce` command on the affected server.

IMPORTANT Only front-end virtual directories are listed in Exchange Admin Center and only the settings of front-end virtual directories are modified by the reset. If you also want to reset the corresponding back-end virtual directory after resetting a front-end virtual directory, you must do this in Exchange Management Shell.

In Exchange Management Shell, you have additional management options for the Autodiscover service. To get detailed information about the Autodiscover configuration, type the following command:

```
Get-AutodiscoverVirtualDirectory-Server MyServer | fl
```

where *MyServer* is the name of the Mailbox server you want to examine. Included in the detailed information is the identity of the Autodiscover virtual directory, which you can use with related cmdlets, and the authentication methods enabled for internal and external access. By default, Autodiscover is configured to use Basic authentication, NTLM authentication, integrated Windows authentication, Web Services security and Outlook Authorization Authentication. By using the `Set-AutodiscoverVirtualDirectory` cmdlet, you can enable or disable these authentication methods, as well as digest authentication. You can also set the internal and external URLs for Autodiscover. Neither URL is set by default.

By default, only information about the related front-end virtual directories is included. To add information about the related back-end virtual directories, you need to set `-ShowMailboxVirtualDirectories` to `$true`. Set `-ADPropertiesOnly` to `$true` if you want to only view the properties stored in Active Directory. The following example gets information for all Autodiscover virtual directories in the Exchange organization:

Get-AutodiscoverVirtualDirectory -ShowMailboxVirtualDirectories | fl

To disable Autodiscover, type the following command:

```
Remove-AutodiscoverVirtualDirectory -Identity ServerName\DirName  
(WebSiteName)
```

where *ServerName* is the name of the Mailbox server on which this feature should be disabled, *DirName* is the name of the virtual directory to remove, and *WebSiteName* is the name of the web site you are configuring, such as:

```
Remove-AutodiscoverVirtualDirectory -Identity  
"CorpMailSvr25\Autodiscover (Default Web Site)"
```

If you later want to enable Autodiscover, you can type the following command:

```
New-AutodiscoverVirtualDirectory -Identity -Identity "CorpMailSvr25\Autodiscover (Default  
Web Site)"
```

where *MyServer* is the name of the Mailbox server on which this feature should be enabled for the Default Web Site.

Listings 27-1 to 27-4 provide the full syntax and usage for the Get-AutodiscoverVirtualDirectory, New-AutodiscoverVirtualDirectory, Set-AutodiscoverVirtualDirectory and Remove-AutodiscoverVirtualDirectory cmdlets, respectively.

LISTING 27-1 Get-AutodiscoverVirtualDirectory cmdlet syntax and usage

Syntax

```
Get-AutodiscoverVirtualDirectory [-Server ServerName | -Identity VirtualDirID]  
[-ADPropertiesOnly <$true | $false>] [-DomainController DCName]  
[-ShowMailboxVirtualDirectories <$true | $false>]
```

Usage

```
Get-AutodiscoverVirtualDirectory  
-Identity "CorpMailSvr25\Autodiscover (Default Web Site)"
```

LISTING 27-2 New-AutodiscoverVirtualDirectory cmdlet syntax and usage

Syntax

```
New-AutodiscoverVirtualDirectory [-ApplicationRoot RootPath]  
[-AppPoolId AppPoolIdentity] [-BasicAuthentication <$true | $false>]  
[-DigestAuthentication <$true | $false>] [-DomainController DCName]  
[-ExternalURL ExternalURL] [-InternalURL InternalURL]  
[-OAuthAuthentication <$true | $false>]  
[-Path FileSystemPath] [-Role <ClientAccess | Mailbox>]  
[-Server ServerName] [-WebSiteName WebSiteName]
```

[-WindowsAuthentication <\$true | \$false>]

[-WSSecurityAuthentication <\$true | \$false>]

Usage

New-AutodiscoverVirtualDirectory -WebSiteName "Default Web Site"

-BasicAuthentication \$true –WindowsAuthentication \$true

-OAuthAuthentication \$true –WSSecurityAuthentication \$true

New-AutodiscoverVirtualDirectory -WebSiteName "Exchange Back End"

-BasicAuthentication \$true –WindowsAuthentication \$true

-OAuthAuthentication \$true –WSSecurityAuthentication \$true

-Role Mailbox

LISTING 27-3 Set-AutodiscoverVirtualDirectory cmdlet syntax and usage

Syntax

Set-AutodiscoverVirtualDirectory -Identity DirectoryIdentity

[-BasicAuthentication <\$true | \$false>]

[-DigestAuthentication <\$true | \$false>]

[-DomainController DCName]

[-ExternalURL ExternalURL] [-InternalURL InternalURL]

[-LiveIdBasicAuthentication <\$true | \$false>]

[-LiveIdNegotiateAuthentication <\$true | \$false>]

[-OAuthAuthentication <\$true | \$false>]

[-WindowsAuthentication <\$true | \$false>]

[-WSSecurityAuthentication <\$true | \$false>]

Usage

Set-AutodiscoverVirtualDirectory

-Identity "CorpMailSvr25\Autodiscover(Default Web Site)"

-BasicAuthentication \$false -DigestAuthentication \$false

–WindowsAuthentication \$true

LISTING 27-4 Remove-AutodiscoverVirtualDirectory cmdlet syntax and usage

Syntax

Remove-AutodiscoverVirtualDirectory -Identity DirectoryIdentity

[-DomainController DCName]

Usage

Remove-AutodiscoverVirtualDirectory

-Identity "CorpMailSvr25\Autodiscover (Default Web Site)"

Using Direct Push

Direct Push automates the synchronization process, enabling a mobile device to make requests to keep itself up to date. When the website used with Exchange ActiveSync has SSL enabled, Direct Push allows a mobile device to issue long-lived Hypertext Transfer Protocol Secure (HTTPS) monitoring requests to Exchange Server. Exchange Server monitors activity in the related user's mailbox. If new mail arrives or other changes are made to the mailbox—such as modifications to calendar or contact items—Exchange sends a response to the mobile device, stating that changes have occurred and that the device should initiate synchronization with Exchange Server. The device then issues a synchronization request. When synchronization is complete, the device issues another long-lived HTTPS monitoring request.

Port 443 is the default TCP port used with SSL. For Direct Push to work, port 443 must be opened between the Internet and the organization's Internet-facing Mailbox server or servers. You do not need to open port 443 on your external firewalls to all of your Mailbox servers—only those to which users can establish connections. The Mailbox server receiving the request automatically proxies the request so that it can be handled appropriately. If necessary, this can also mean proxying requests between the mobile device and the Mailbox server in the user's home site. A user's home site is the Active Directory site in which the Mailbox server hosting his or her mailbox is located.

TIP On your firewall, Microsoft recommends increasing the maximum time-out for connections to 30 minutes to help optimize the efficiency of direct push.

Using Remote Device Wipe

Although passwords help to protect mobile devices, they don't prevent access to the device. You can protect the data on mobile devices in several ways. One such way is to apply a mobile device mailbox policy that controls access to the device and encrypts its content. Another way is to have a strict policy that requires users and administrators to remotely wipe lost or stolen devices. A remote device wipe command instructs a mobile device to delete all Exchange and corporate data.


Remotely Wiping a Device

An administrator or the owner of the device can prevent the compromising of sensitive data by initiating a remote device wipe. After you initiate a remote device wipe and the device receives the request, the device confirms the remote wipe request by sending a confirmation message and then removes all its sensitive data the next time it connects to Exchange Server. Wiping sensitive data should prevent it from being compromised.

The way remote wipe is implemented depends on the way the related protocol is implemented on the device. Although Exchange 2016 only requires that Exchange and corporate data be removed, most device operating systems wipe all data on the device and then return the device to its factory default condition. A complete wipe can also remove any data stored on any storage card inserted into the device. On the other hand, when you issue a remote wipe for a device that fully support Exchange, the wipe typically only affects Exchange and corporate data. For these devices, client application settings also can determine whether the wipe actually deletes the sensitive data or simply makes it inaccessible. As data on these devices is encrypted by default, any data remaining would be protected by encryption.

The easiest way to wipe a device remotely is to have the device owner initiate the wipe using Outlook Web App. When the device acknowledges the request, the user will get a confirmation email.

The device owner can wipe a device by following these steps:

1. Open your web browser. In the Address field, type the Outlook Web App URL, such as <https://mail.tvpress.com/owa>, and then press Enter to access this page.
2. When prompted, provide the logon credentials of the user whose device you want to wipe. Do not provide your administrator credentials.
3. On the Outlook Web App toolbar, click Settings () , and then click Options.
4. The left pane of the Options view provides a list of options. Click the General heading to expand it and then select Mobile Devices.
5. The user's mobile devices are listed in the details pane. Select the device you want to wipe, and then click Wipe Device.

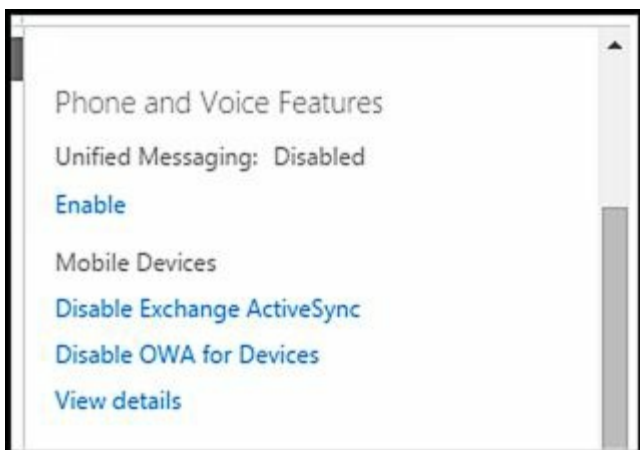
6. Confirm the action when prompted.
7. Track the status of the device. When the status changes from Wipe Pending to Wipe Successful, the device wipe is complete.

NOTE You can use Outlook Web App for remote device wiping only if the user has used the device previously to access Exchange Server and if you have enabled the Segmentation feature of Exchange Active Directory Integration (which is the default configuration).

CAUTION Because wiping a device causes complete data loss, you should do this only when you've contacted the user directly (preferably in person) and confirmed that the mobile device has been lost and that he or she understands the consequences of wiping the device. If your organization has a formal policy regarding the wiping of lost devices that might contain sensitive company data, be sure you follow this policy and get any necessary approvals. Keep in mind that although a remote wipe makes it very difficult to retrieve any data from the device, in theory retrieval is possible with sophisticated data recovery tools.

Alternatively, an administrator can log on to Exchange Admin Center and initiate a remote wipe by completing the following steps:

1. In Exchange Admin Center, select Recipients in the Navigation menu, and then select Mailboxes.
2. Select the mailbox for the user whose device you want to wipe. Next, in the details pane, under Mobile Devices, click View Details.



3. On the Mobile Device Details page, select the lost device, and then select Wipe Data.
4. Click Save to initiate the remote wipe.
5. Track the status of the device. When the status changes from Wipe Pending to Wipe Successful, the device wipe is complete.

In Exchange Management Shell, you can examine and filter through all of the mobile devices that have linked to Exchange by using Get-MobileDevice. You also can list the mobile devices registered as partners for a user's mailbox by using the Get-MobileDeviceStatistics cmdlet. In either case, the device identity you want is the DeviceId string. If the user has multiple mobile devices, also be sure to consult the

DeviceModel and DeviceOperatorNetwork values.

After you know the mobile device identity, you can issue a remote device wipe command by using the Clear-MobileDevice cmdlet. You then need to confirm that you want to wipe the device when prompted by pressing the Y key. Listings 27-5 to 27-7 provide the syntax and usage for Get-MobileDevice, Get-MobileDeviceStatistics, and Clear-MobileDevice cmdlets, respectively. With Get-MobileDeviceStatistics, you can specify either the unique identity of the remote device or the user mailbox with which you want to work. The –GetMailboxLog parameter retrieves mailbox logs and usage information. Use the –OutputPath parameter to direct the statistics to a specific folder path or the –NotificationEmailAddresses parameter to email the statistics to specified email addresses.

IMPORTANT If you determine that you’ve made a mistake in issuing a remote wipe, you should immediately issue a cancellation request by using the Clear-MobileDevice cmdlet. In this case, set the –Cancel parameter to \$true. The remote device processes the cancellation request only if the remote wipe has not yet been initiated.

NOTE Exchange also supports the Get-ActiveSyncDevice, Get-ActiveSyncDeviceStatistics and Clear-ActiveSyncDevice cmdlets, which have similar syntax and options as Get-MobileDevice, Get-MobileDeviceStatistics, and Clear-MobileDevice respectively. As the ActiveSyncDevice cmdlets only work with ActiveSync devices and the MobileDevice cmdlets work with all supported devices, I prefer to use the MobileDevice cmdlets and you probably will too.

LISTING 27-5 Get-MobileDevice cmdlet syntax and usage

Syntax

```
Get-MobileDevice [-Identity MobileDeviceId] {AddtlParams}
Get-MobileDevice -Mailbox MailboxId {AddtlParams}
{AddtlParams}
[-ActiveSync <$true | $false>] [-DomainController FullyQualifiedName]
[-Filter FilterValues] [-Monitoring <$true | $false>]
[-Organization OrgId] [-OrganizationalUnit OUIId]
[-OWAforDevices <$true | $false>] [-ResultSize Size]
[-SortBy AttributeName]
```

Usage

```
Get-MobileDevice -OrganizationalUnit Sales
```

LISTING 27-6 Get-MobileDeviceStatistics cmdlet syntax and usage

Syntax

```
Get-MobileDeviceStatistics -Identity MobileDeviceId {AddtlParams}
Get-MobileDeviceStatistics -Mailbox MailboxId {AddtlParams}
{AddtlParams}
[-ActiveSync <$true | $false>] [-DomainController FullyQualifiedName]
[-GetMailboxLog <$true | $false>] [-NotificationEmailAddresses
email1,email2,...emailN] [-OWAMobileApp <$true | $false>]
[-ShowRecoveryPassword <$true | $false>]
```

Usage

```
Get-MobileDeviceStatistics -Mailbox "David Pelton"
```

LISTING 27-7 Clear-MobileDevice cmdlet syntax and usage

Syntax

```
Clear-MobileDevice -Identity MobileDeviceId
[-Cancel <$true | $false>] [-DomainController FullyQualifiedName]
[-NotificationEmailAddresses email1,email2,...emailN]
```

Usage

```
Clear-MobileDevice -Identity "Mobile_DavidP"
Clear-MobileDevice -Identity "Mobile_DavidP" -Cancel $true
```

Reviewing the Remote Wipe Status

When you initiate a remote wipe, the mobile device removes all its data the next time it connects to Exchange Server. You can review the remote wipe status by using an alternate syntax for the Get-MobileDeviceStatistics cmdlet. Instead of passing the –Mailbox parameter to the cmdlet, use the –Identity parameter to specify the DeviceId string of the device you wiped. The statistics returned will include these output parameters:

- **DeviceWipeRequestTime** The time you requested a remote wipe
- **DeviceWipeSentTime** The time the server sent the remote wipe command to the device
- **DeviceWipeAckTime** The time when the device acknowledged receipt of the remote wipe command


If there is a DeviceWipeSentTime timestamp, the device has connected to Exchange Server and Exchange Server sent the device the remote wipe command. If there is a DeviceWipeAckTime timestamp, the device acknowledged receipt of the remote wipe and has started to wipe its data.

Using Password Recovery

Users can create passwords for their mobile devices. If a user forgets his password, you can obtain a recovery password that unlocks the device and lets the user create a new password. The user can also recover his device password by using Outlook Web App.

Recovering a Device Password

To use Outlook Web App to recover a user's device password, complete the following steps:

1. Open a web browser. In the Address field, type the Outlook Web App URL, such as <https://mail.tvpress.com/owa>, and then press Enter to access this page.
2. When prompted, have the user enter her logon credentials or provide the user's logon credentials. Do not provide your administrator credentials.
3. On the Outlook Web App toolbar, click Settings () , and then click Options.
4. The left pane of the Options view provides a list of options. Click the General heading to expand it and then select Mobile Devices.
5. The user's mobile devices are listed in the details pane. Select the device for which you are recovering the password.
6. Click Display Recovery Password.

You also can display the device recovery password by completing the following steps:

1. In the Exchange Admin Center, select Recipients in the Navigation menu, and then select Mailboxes.
2. Select the mailbox for the user whose device you want to wipe. Next, in the details pane, under Mobile Devices, click View Details.
3. The device recovery password is listed.

In Exchange Management Shell, you can display the device recovery password by using the `-ShowRecoveryPassword` parameter of the `Get-MobileDeviceStatistics` cmdlet.

Listing 27-8 provides the syntax and usage.

LISTING 27-8 Recovering a device password

Syntax

```
Get-MobileDeviceStatistics -Mailbox MailboxIdentity  
-ShowRecoveryPassword $true {AddtlParams}
```

```
Get-MobileDeviceStatistics -Identity MobileDeviceIdentity
```

-ShowRecoveryPassword \$true {AddtlParams}

{AddtlParams}

[-ActiveSync <\$true | \$false>] [-DomainController FullyQualifiedName]

[-GetMailboxLog <\$true | \$false>] [-NotificationEmailAddresses

email1,email2,...emailN] [-OWAforDevices <\$true | \$false>]

Usage

Get-MobileDeviceStatistics -Mailbox "HelenB@tvpress.com"

-ShowRecoveryPassword \$true

Managing File Access and Document Viewing

Exchange 2016 includes many features designed to make it easier for users to work with files and documents. Users can access files directly, remotely and via Office Web Apps Server viewing. Each of which has separate configuration options.

Configuring Direct File Access

Exchange Server 2016 allows users to access files directly through Outlook, Outlook Web App, and related services by default. This means that users will be able to access files attached to email messages. You can configure how users interact with files direct file access by using one of three options in the Exchange Admin Center:

- **Allow** Allows users to access files of the specified types, and sends the users' browser information that allows the files to be displayed or opened in the proper applications
- **Block** Prevents users from accessing files of the specified types
- **Force Save** Forces users to save files of the specified types prior to opening them

In a standard configuration, Exchange 2016 allows, blocks and force saves many file extensions and Multipurpose Internet Mail Extensions (MIME) values. Allowed files include .avi, .bmp, .doc, .docm, .docx, .gif, .jpg, .mp3, and other standard file types. Blocked files include that contain executables and scripts, such as .bat, .cmd,.exe, .ps1, .vbe, .vbs, and text/javascript. Forced save files include specific types of application files, such as .dcr, .dir, .spl, and .swf.

NOTE The related settings are applied to the Outlook Web App (OWA) virtual directory on Mailbox servers. If a server has multiple OWA virtual directories or you have multiple Mailbox servers, you must configure each directory and server separately. If there are conflicts between the allow, block, and force save lists, the allow list takes precedence, which means that the allow list settings override the block list and the force save list. As updates are applied to Exchange Server, the default lists can change. Be sure to check the currently applied defaults.

Exchange Server considers all file extensions and MIME types not listed on the allow, block, or force save list to be unknown. The default setting for unknown file types is force save.


Based on the user's selection, the configuration of her network settings, or both, Exchange divides all client connections into one of two classes:

- **Public or shared computer** A public computer is a computer being used on a public network or a computer shared by multiple people.
- **Private computer** A private computer is a computer on a private network that is used by one person.

For each Mailbox server, you can enable or disable direct access to files separately for

public computers and private computers. However, the allow, block, and force save settings for both types of computers are shared and applied to both public and private computers in the same way.

You can configure direct file access on front-end virtual directories by completing the following steps:

1. In the Exchange Admin Center, select Servers in the Navigation menu, and then select Virtual Directories to view a list of the front-end virtual directories used by Mailbox servers in the Exchange organization.
2. Select the OWA virtual directory on the Mailbox server you want to manage, and then select Edit (). Typically, you'll want to configure the OWA virtual directory on the Default Web Site because this directory is used by default for Outlook Web App.
3. In the Virtual Directory dialog box, select the File Access page.
4. To enable or disable direct file access for public computers, under Public Or Shared Computer, select or clear the Direct File Access check box, as appropriate.



IMPORTANT When you disable features in the front end, you prevent them from being used because the front end proxies connections to the back end and blocks disabled features from being used. However, if you enable a feature in the front end but the feature is disabled in the back end, clients also won't be able to use the feature.

5. Under Private Computer, you can select or clear the Direct File Access check box to enable or disable direct file access for private computers.
6. Click Save to apply your settings. As necessary, make corresponding changes in the related back-end virtual directory using Exchange Management Shell.

In Exchange Management Shell, you can use the Set-OWAVirtualDirectory cmdlet to manage the direct file-access configuration. Use the `-Identity` parameter to identify the virtual directory with which you want to work, such as:

```
Set-OWAVirtualDirectory -Identity "Corpsvr127\owa (Default Web Site)"
```

Then specify how you want to configure direct file access on the front-end and back-end virtual directory, such as:

```
Set-OWAVirtualDirectory -Identity "Corpsvr127\owa (Default Web Site)"  
-DirectFileAccessOnPublicComputersEnabled $false  
-DirectFileAccessOnPrivateComputersEnabled $true
```

```
Set-OWAVirtualDirectory -Identity "Corpsvr127\owa (Exchange Back End)"  
-DirectFileAccessOnPublicComputersEnabled $false  
-DirectFileAccessOnPrivateComputersEnabled $true
```

If you are unsure of the virtual directory identity value, use the `Get-OWAVirtualDirectory` cmdlet to retrieve a list of available virtual directories on a named server, as shown in the following example:

```
Get-OWAVirtualDirectory -Server "Corpsvr127" -ShowMailboxVirtualDirectories
```

Alternatively, you could get the `OWAVirtualDirectory` object for both the front end and back end and then set the desired options on both as shown in the following example:

```
Get-OWAVirtualDirectory -Server Corpsvr127 -ShowMailboxVirtualDirectories |  
Set-OWAVirtualDirectory -DirectFileAccessOnPublicComputersEnabled $false  
-DirectFileAccessOnPrivateComputersEnabled $true
```

You could just as easily apply the changes to multiple Exchange servers throughout the organization. If you want to make changes across all servers, however, I recommend adding the `-Whatif` parameter to ensure the command is going to work exactly as expected before executing the command without the `-Whatif` parameter. In the following example, you disable direct file access on public computers on all front-end and back-end OWA virtual directories:

```
Get-OWAVirtualDirectory -ShowMailboxVirtualDirectories |  
Set-OWAVirtualDirectory -DirectFileAccessOnPublicComputersEnabled $false  
-WhatIf
```

You configure allowed file types and allowed MIME types by using the `-AllowedFileTypes` and `-AllowedMIMETypes` parameters respectively. As these are multivalued properties, you must either enter the complete set of allowed values or use a special shorthand to insert into or remove values from these multivalued properties. The shorthand for adding values is:

```
@{Add="<ValuetoAdd1>","<ValuetoAdd2>"...}
```

As you'll typically want to configure the front-end and back-end virtual directories in the same way, the following example sets the allowed file types on both the front-end and back-end OWA virtual directories:

```
Get-OWAVirtualDirectory -Server Corpsvr127 -ShowMailboxVirtualDirectories |  
Set-OWAVirtualDirectory -AllowedFileTypes @{Add=".log",".man"}
```

The shorthand for removing values is:

```
@{Remove="<ValuetoRemove1>","<ValuetoRemove2>"...}
```

Such as:

```
Get-OWAVirtualDirectory -Server Corpsvr127 -ShowMailboxVirtualDirectories |  
Set-OWAVirtualDirectory -AllowedFileTypes @{Remove=".log",".man"}
```

If you want to add values and remove others, you can do this as well by using the following shorthand:

```
@{Add="<ValuetoAdd1>","<ValuetoAdd2>"...;  
Remove="<ValuetoRemove1>","<ValuetoRemove1>"...}
```

You can confirm that values were added or removed as expected by using `Get-OWAVirtualDirectory`. Because there are so many allowed file types, you won't get a complete list of file types if you examine the `-AllowedFileTypes` property as shown in the following example:

```
Get-OWAVirtualDirectory -Identity "Corpsvr127\owa (Default Web Site)" |  
fl name, allowedfiletypes
```

A workaround to examine all the values of such a property follows:

```
$vdir = Get-OWAVirtualDirectory -Identity "Corpsvr127\owa (Default Web  
Site)"
```

```
$data = $vdir.allowedfiletypes  
$data | fl *
```

In this case, you store the virtual directory object in the `$vdir` variable. Next, you store the values associated with this object's `AllowedFileTypes` parameter in the `$data` variable. Finally, you list each allowed file type.

You can use similar techniques to work with

- **Blocked file types and blocked MIME types** The corresponding parameters are `-BlockedFileTypes` and `-BlockedMimeTypes` respectively.
- **Forced Save file types and forced save MIME types** The corresponding parameters are `-ForcedSaveFileTypes` and `-ForcedSaveMimeTypes` respectively.

Configuring Remote File Access

Exchange Server 2016 allows users to access files remotely through Outlook Web App (OWA) by default. This means users will be able to access Windows SharePoint Services and Universal Naming Convention (UNC) file shares on SharePoint sites. SharePoint sites consist of Web Parts and Windows ASP.NET–based components that allow users to share documents, tasks, contacts, events, and other information. When you configure UNC file shares on SharePoint sites, you enable users to share folders and files.

You configure remote file access by using configuration options for the ActiveSync virtual directory. The `-RemoteDocumentsBlockedServers` and -

RemoteDocumentsAllowedServers parameters of the Set-ActiveSyncVirtualDirectory cmdlet specify the host names of servers from which clients are denied or allowed access respectively. If there is a conflict between the blocked servers list and the allowed servers list, the block list takes precedence.

As the -RemoteDocumentsBlockedServers and -RemoteDocumentsAllowedServers parameters are multivalued properties, you must either enter the complete set of allowed values or use the special shorthand discussed earlier in this chapter in the "Configuring Direct File Access" section to insert into or remove values from these multivalued properties. To add a server to the blocked or allowed servers list, use the fully qualified domain name of the server, such as *mailsvr83.tvpress.com*.

The following example adds two servers to the allowed servers list throughout the Exchange organization:

```
Get-ActiveSyncVirtualDirectory -ShowMailboxVirtualDirectories |  
Set-OWAVirtualDirectory -RemoteDocumentsAllowedServers  
@{Add="mailsvr83.tvpress.com","corpserver18.treyresearch.net"}
```

Servers that are not listed on either the allow list or the block list are considered to be unknown servers. By default, access to unknown servers is allowed. You can use the -RemoteDocumentsActionForUnknownServers parameter to specify whether to allow or block unknown servers. Set the parameter value to Allow or Block as appropriate. Here is an example:

```
Get-ActiveSyncVirtualDirectory -ShowMailboxVirtualDirectories |  
Set-OWAVirtualDirectory -RemoteDocumentsActionForUnknownServers Block
```

Users have access only to shares hosted on internal servers. For a server to be considered an internal server, you must tell Exchange about the domain suffixes that should be handled as internal by using the -RemoteDocumentsInternalDomainSuffixList parameter. This is a multivalued parameter.

To add a domain suffix, specify the fully qualified domain name of the suffix. An example follows:

```
Get-ActiveSyncVirtualDirectory -ShowMailboxVirtualDirectories |  
Set-OWAVirtualDirectory -RemoteDocumentsInternalDomainSuffixList  
@{Add="tvpress.com","treyresearch.net"}
```

To remove a domain suffix, specify the suffix to remove, such as:

```
Get-ActiveSyncVirtualDirectory -ShowMailboxVirtualDirectories |  
Set-OWAVirtualDirectory -RemoteDocumentsInternalDomainSuffixList  
@{Remove="proseware.com","litwareinc.com"}
```

Integrating Office Web Apps Servers

Although earlier releases of Exchange included functionality for viewing documents directly in Outlook Web App, Microsoft has since developed separate server functionality to provide full viewing and editing functionality. The new architecture

requires installing Office Web Apps servers which are then integrated into your Exchange organization to provide viewing and editing functions for Office documents.

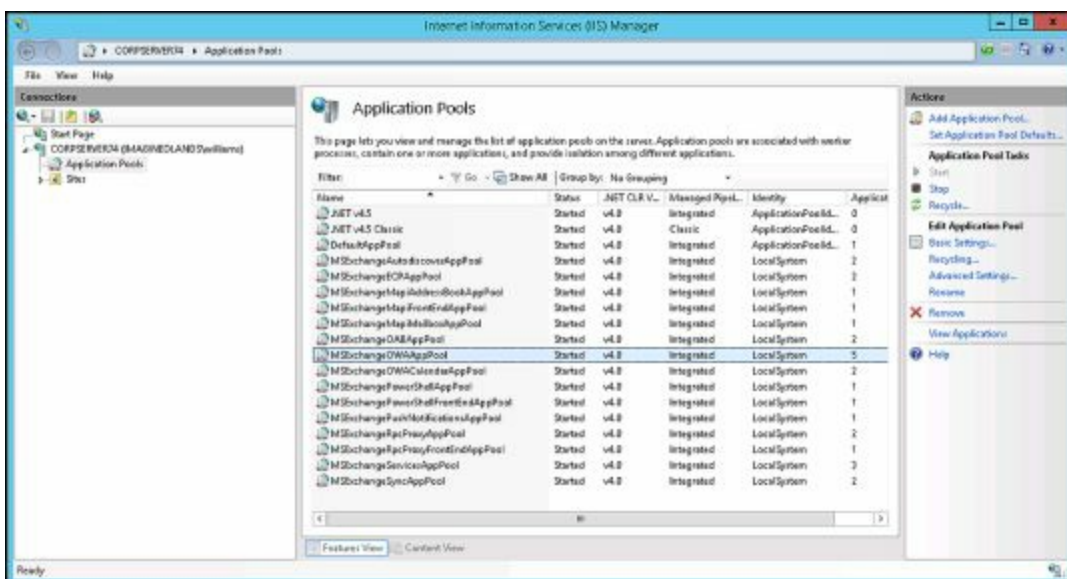
After you install Office Web Apps servers on your network, you need to perform a series of steps to prepare your Mailbox servers to use the new architecture:

1. Use the `-WacDiscoveryEndpoint` parameter of the `Set-OrganizationConfig` cmdlet to specify the Discovery URL for Office Web Apps servers, such as:

`Set-OrganizationConfig -WacDiscoveryEndpoint`

`https://MailServer85/hosting/discovery`

2. Log on to the Mailbox server. Start IIS Manager. In Server Manager, click Tools, and then select Internet Information Services (IIS) Manager.
3. In IIS Manager, expand the base node for the server you want to work with and then select Application Pools.



4. Click the `MSExchangeOWAAppPool` to select the application pool for OWA.
5. In the Actions pane, click `Recycle` to recycle the application pool. When the application pool restarts, OWA will detect the Discovery endpoint and then your Office Web Apps servers will be used for viewing and editing Office documents.

You can confirm integration and availability of the Discovery endpoint by entering the following command:

`Get-OrganizationConfig | fl WacDiscoveryEndpoint`

The output should include the URL you previously specified.

Once you've installed Office Web Apps servers and used the previous procedure to integrate them into your Exchange organization, OWA users will be able to view and edit Office documents without having the applications associated with those file types installed on their computing devices. Files types that WAC Viewing allows users to view and edit include:

- Microsoft Office Excel spreadsheets with the `.xls`, `.xlsx`, `.xlm`, and `.xlsb` extensions
- Microsoft Office Word documents with the `.doc`, `.docx`, `.dot`, `.dotx`, and `.dotm`

extensions

- Microsoft Office PowerPoint presentations with the .pps, .ppsx, .ppt, .pptx, .pot, .potm, and .ppsm extensions

When there are conflicting settings between the direct file, remote file, and WAC Viewing settings, you can force clients to use WAC Viewing first, if you want. This ensures that the documents will be opened using Office Web Apps servers as helpers.

You can enable or disable WAC Viewing separately for public computers and private computers. However, supported document settings for both types of computers are shared and applied to both public and private computers in the same way.

In Exchange Management Shell, you can use the Set-OWAVirtualDirectory cmdlet to manage the WAC Viewing configuration. Use the –Identity parameter to identify the virtual directory you want to work, such as:

```
Set-OWAVirtualDirectory –Identity "Corpsvr127\owa (Default Web Site)"
```

Then specify how you want to configure WAC Viewing on the front-end and back-end virtual directory, such as:

```
Set-OWAVirtualDirectory –Identity "Corpsvr127\owa (Default Web Site)"  
-WacViewingOnPublicComputersEnabled $false  
-WacViewingOnPrivateComputersEnabled $true  
-WacEditingEnabled $true
```

```
Set-OWAVirtualDirectory –Identity "Corpsvr127\owa (Exchange Back End)"  
-WacViewingOnPublicComputersEnabled $false  
-WacViewingOnPrivateComputersEnabled $true  
-WacEditingEnabled $true
```

If you are unsure of the virtual directory identity value, use the Get-OWAVirtualDirectory cmdlet to retrieve a list of available virtual directories on a named server, as shown in the following example:

```
Get-OWAVirtualDirectory –Server "Corpsvr127" -ShowMailboxVirtualDirectories
```

Typically, you'll want to configure the front-end and back-end virtual directories in the same way.

Working with Mobile Devices and Device Policies

Mobile device mailbox policy makes it possible to enhance the security of mobile devices used to access your Exchange servers. For example, you can use policy to require a password of a specific length and to configure devices to automatically prompt for a password after a period of inactivity.

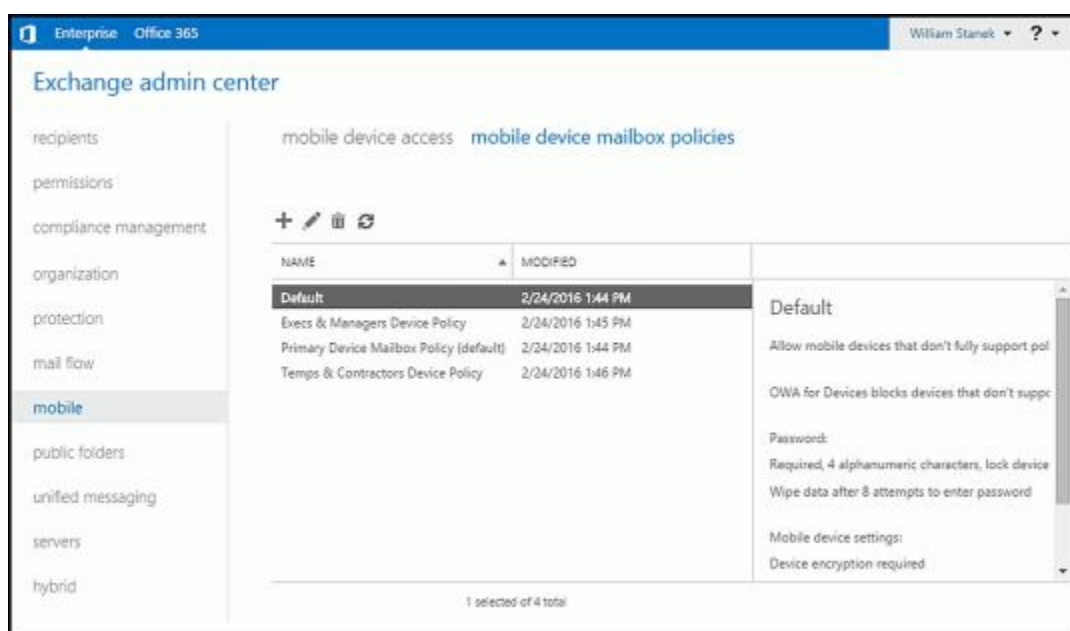
Each mailbox policy you create has a name and a specific set of rules with which it is associated. Because you can apply policies separately to mailboxes when you create or modify them, you can create different policies for different groups of users. For example, you can have one policy for users and another policy for managers. You can also create separate policies for departments within the organization. For example, you can have separate policies for Marketing, Customer Support, and Technology.

NOTE Mobile device mailbox policies replace ActiveSync mailbox policies. If your organization is still using ActiveSync mailbox policies, which are being phased out, you should transition to mobile device policies.

Viewing Existing Mobile Device Mailbox Policies

When the Mailbox role is installed on an Exchange server, the setup process creates a default mobile device mailbox policy, which allows enterprise mobile devices to be used without restrictions or password requirements. All users with mailboxes have this policy applied by default. You can modify the settings of this policy to change the settings for all users or create new policies for specific groups of users.

In the Exchange Admin Center, you can view the currently configured mobile device mailbox policies by selecting Mobile in the Navigation menu, and then selecting Mobile Device Mailbox Policies. In the details pane, you'll see a list of current policies.



In Exchange Management Shell, you can list policies by using the `Get-MobileDeviceMailboxPolicy` cmdlet. Listing 27-9 provides the syntax, usage, and

sample output. If you do not provide an identity with this cmdlet, all available mobile device mailbox policies are listed.

LISTING 27-9 Get-MobileDeviceMailboxPolicy cmdlet syntax and usage

Syntax

```
Get-MobileDeviceMailboxPolicy [-Identity MailboxPolicyId]
[-DomainController FullyQualifiedName] [-Organization OrgId]
```

Usage

```
Get-MobileDeviceMailboxPolicy
Get-MobileDeviceMailboxPolicy -Identity "Primary Device Mailbox Policy"
```

Output


```
RunspaceId           :
AllowNonProvisionableDevices      : True
AlphanumericPasswordRequired      : True
AttachmentsEnabled                : True
DeviceEncryptionEnabled           : False
RequireStorageCardEncryption      : False
PasswordEnabled                   : True
PasswordRecoveryEnabled           : False
DevicePolicyRefreshInterval       : Unlimited
AllowSimplePassword               : False
MaxAttachmentSize                 : Unlimited
WSSAccessEnabled                  : True
UNCAccessEnabled                  : True
MinPasswordLength                 : 4
MaxInactivityTimeLock             : 00:15:00
MaxPasswordFailedAttempts         : 8
PasswordExpiration                 : 90.00:00:00
PasswordHistory                   : 5
IsDefault                         : True
AllowApplePushNotifications       : True
AllowMicrosoftPushNotifications  : True
AllowGooglePushNotifications      : True
AllowStorageCard                  : True
AllowCamera                       : True
RequireDeviceEncryption           : True
AllowUnsignedApplications         : True
AllowUnsignedInstallationPackages : True
AllowWiFi                         : True
AllowTextMessaging                : True
AllowPOPIMAPEmail                 : True
AllowIrDA                         : True
RequireManualSyncWhenRoaming      : False
AllowDesktopSync                  : True
AllowHTMLEmail                    : True
RequireSignedSMIMEMessages        : False
RequireEncryptedSMIMEMessages     : False
```

AllowSMIMESoftCerts : True
 AllowBrowser : True
 AllowConsumerEmail : True
 AllowRemoteDesktop : True
 AllowInternetSharing : True
 AllowBluetooth : Allow
 MaxCalendarAgeFilter : All
 MaxEmailAgeFilter : All
 RequireSignedSMIMEAlgorithm : SHA1
 RequireEncryptionSMIMEAlgorithm : TripleDES
 AllowSMIMEEncryptionAlgorithmNegotiation : AllowAnyAlgorithmNegotiation
 MinPasswordComplexCharacters : 3
 MaxEmailBodyTruncationSize : Unlimited
 MaxEmailHTMLBodyTruncationSize : Unlimited
 UnapprovedInROMApplicationList : {}
 ApprovedApplicationList : {}
 AllowExternalDeviceManagement : False
 MobileOTAUpdateMode : MinorVersionUpdates
 AllowMobileOTAUpdate : True
 IrmEnabled : True
 AdminDisplayName :
 ExchangeVersion : 0.1 (8.0.535.0)
 Name : Default
 DistinguishedName : CN=Default,CN=Mobile Mailbox”
 Policies,CN=First Organization,CN=MicrosoftExchange,
 CN=Services,CN=Configuration,DC=imaginedlands,DC=local
 Identity : Default
 ...
 Id : Default
 OriginatingServer : CorpServer91.imaginedlands.local
 IsValid : True
 ObjectState : Unchanged

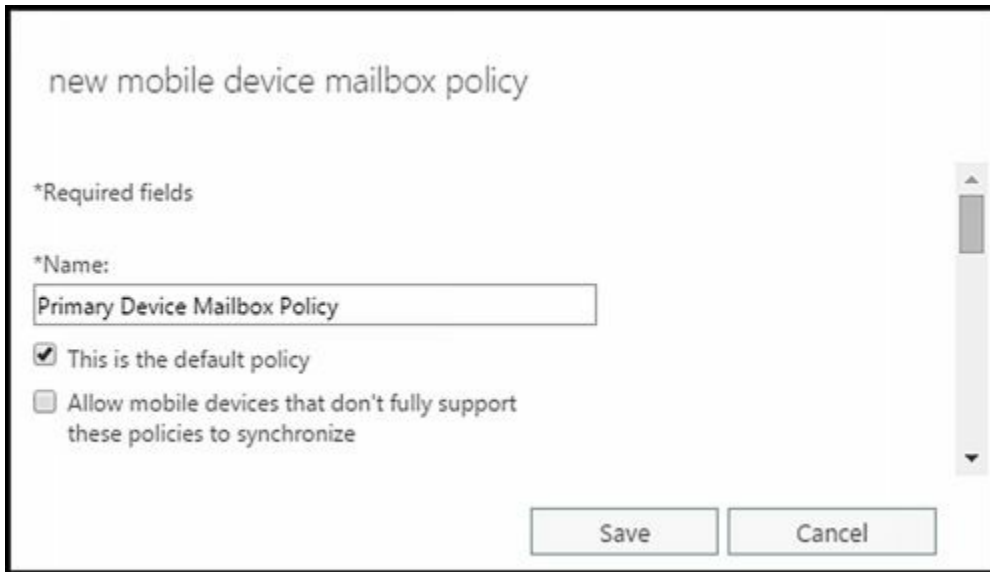
Creating Mobile Device Mailbox Policies

The mobile device mailbox policies you create apply to your entire organization. You apply policies separately after you create them, as discussed later in this chapter in the “Assigning Mobile Device Mailbox Policies” section.

You can create a new policy by completing the following steps:

1. In Exchange Admin Center, select **Mobile** in the Navigation menu, and then select **Mobile Device Mailbox Policies** to see a list of currently defined mobile device mailbox policies.
2. Click **New** () to open the **New Mobile Device Mailbox Policy** dialog box.
3. Type a descriptive name for the policy. If you want the policy to be assigned to all users who are currently using the previously assigned default policy, select **This Is The Default Policy**.
4. By default, mobile devices that do not support all device mailbox settings can't

synchronize with Exchange. This is by design to ensure strict security can be enforced. If you want to allow older devices to sync with Exchange regardless of whether they fully support device policy, select the Allow Mobile Devices That Don't Fully Support... checkbox.



new mobile device mailbox policy

*Required fields

*Name:
Primary Device Mailbox Policy

This is the default policy

Allow mobile devices that don't fully support these policies to synchronize

Save Cancel

5. Before you can apply policy restrictions, you must specify that device passwords are required by selecting the Require A Password checkbox. If you do not select this option, you cannot specify password requirements.
6. Next, use the following options provides to specify the password requirements:
 - **Allow Simple Passwords** Allows the user to use a noncomplex password instead of a password that meets the minimum complexity requirements.
 - **Require An Alphanumeric Password** Requires that a password contain numeric and alphanumeric characters. If you do not select this option, users can use simple passwords, which might not be as secure. If you select this option, you can also specify the number of character sets that are required to be used in passwords. The four character sets are lowercase letters, uppercase letters, numbers, and symbols. You can require from one to four of these character sets to be used in passwords.

new mobile device mailbox policy

Require a password

Allow simple passwords

Require an alphanumeric password

Password must include this many character sets:

3

Require encryption on device

Minimum password length:

4

Number of sign-in failures before device is wiped:

10

When you enable a mobile device password, you can specify a variety of password requirements.

[Learn more](#)

Save Cancel

- **Require Encryption On Device** Requires mobile devices to use encryption. Because encrypted data cannot be accessed without the appropriate password, this option helps to protect the data on the device. If you select this option, Exchange allows devices to download data only if they can use encryption (except when you allow mobile devices that don't fully support mobile device mailbox policy).
- **Minimum Password Length** Allows you to set a minimum password length. You must select the related check box to set the minimum password length, such as eight characters. The longer the password, the more secure it is. A good minimum password length is between 8 and 12 characters, which is sufficient in most cases.
- **Number Of Sign-In Failures Before Device Is Wiped** Allows you to specify the number of login failures before the device is wiped. If you select this option, be sure to set a high enough value so that mobile devices aren't accidentally wiped by users. For example, rather than setting a low value, such as 3, use a higher value, such as 9.
- **Require Sign-In After The Device Has Been Inactive For (Minutes)** Allows you to specify the length of time that a device can go without user input before it locks. If you select this option, be sure to set an interval that allows for normal workflow and isn't disruptive. For example, if high security isn't a requirement, you may want to require users to sign-in after 5 to 7 minutes of inactivity rather than having the device lock itself after 2 to 3 minutes of inactivity.
- **Enforce Password Lifetime (Days)** Allows you to specify the maximum length of time users can keep a password before they have to change it. You can use this option to require users to change their passwords periodically. A good password expiration value is between 30 and 90 days. This period is sufficient to allow use of the password without requiring overly frequent changes.
- **Password Recycle Count** Allows you to specify how frequently old passwords can be reused. You can use this option to discourage users from changing back and forth

between a common set of passwords. To disable this option, set the size of the password history to zero. To enable this option, set the desired size of the password history. A good value is between 3 and 6. This helps to deter users from switching between a small list of common passwords.

7. Click Save to create the policy. Optimize the configuration, as discussed in the following section of this chapter, “Optimizing Mobile Device Mailbox Policies.”

In Exchange Management Shell, you can create new mobile device mailbox policies by using the `New-MobileDeviceMailboxPolicy` cmdlet. Listing 27-10 provides the syntax and usage. There are additional policy settings you can access in the shell that you cannot access in the Exchange Admin Center.

LISTING 27-10 `New-MobileDeviceMailboxPolicy` cmdlet syntax and usage

Syntax

```
New-MobileDeviceMailboxPolicy -Name PolicyName
[-AllowBluetooth <Disable | HandsfreeOnly | Allow>]
[-AllowSMIMEEncryptionAlgorithmNegotiation <BlockNegotiation |
OnlyStrongAlgorithmNegotiation | AllowAnyAlgorithmNegotiation>]
[-ApprovedApplicationList AppList] [-DevicePolicyRefreshInterval
Interval] [-DomainController FullyQualifiedName]
[-MaxAttachmentSize MaxSizeKB] [-MaxCalendarAgeFilter <All | TwoWeeks |
OneMonth | ThreeMonths | SixMonths>] [-MaxEmailAgeFilter <All |
OneDay | ThreeDays | OneWeek | TwoWeeks | OneMonth>]
[-MaxEmailBodyTruncationSize MaxSizeKB]
[-MaxEmailHTMLBodyTruncationSize MaxSizeKB] [-MaxInactivityTimeLock
InactiveTime] [-MaxPasswordFailedAttempts NumAttempts]
[-MinPasswordComplexCharacters MinComplexChars] [-MinPasswordLength
MinLength] [-MobileOTAUpdateMode <MajorVersionUpdates |
MinorVersionUpdates | BetaVersionUpdates>] [-Organization
OrgId] [-PasswordExpiration PasswordExp] [-PasswordHistory HistLength]
[-RequireEncryptionSMIMEAlgorithm <TripleDES | DES | RC2128bit |
RC264bit | RC240bit>] [-RequireSignedSMIMEAlgorithm <SHA1 | MD5>]
[-UnapprovedInROMApplicationList AppList] {OptionalTrueFalseParams}
```

{OptionalTrueFalseParams}

```
-AllowBrowser, -AllowCamera, -AllowConsumerEmail, -AllowDesktopSync,
-AllowExternalDeviceManagement, -AllowHTMLEmail, -AllowInternetSharing,
-AllowIrDA, -AllowMicrosoftPushNotifications, -AllowMobileOTAUpdate,
-AllowNonProvisionableDevices, -AllowPOPIMAPEmail,
-AllowRemoteDesktop, -AllowSimplePassword, -AllowSMIMESoftCerts,
-AllowStorageCard, -AllowTextMessaging, -AllowUnsignedApplications,
-AllowUnsignedInstallationPackages, -AllowWiFi,
-AlphanumericPasswordRequired, -AttachmentsEnabled,
-DeviceEncryptionEnabled, -IrmEnabled, -IsDefault, -PasswordEnabled,
-PasswordRecoveryEnabled, -RequireDeviceEncryption,
-RequireEncryptedSMIMEMessages, -RequireManualSyncWhenRoaming,
-RequireSignedSMIMEMessages, -RequireStorageCardEncryption,
```


-UNCAccessEnabled, -WSSAccessEnabled

Usage

New-MobileDeviceMailboxPolicy -Name "Primary Mobile Device Mailbox Policy"

-AllowNonProvisionableDevices \$true

-PasswordEnabled \$true

-AlphanumericPasswordRequired \$true

-MaxInactivityTimeLock "00.15:00"

-MinPasswordLength "8"

-PasswordRecoveryEnabled \$true

-RequireDeviceEncryption \$true

-AttachmentsEnabled \$true

Optimizing Mobile Device Mailbox Policies

When you create a mobile device mailbox policy, some additional settings are configured automatically. You can modify policy settings by using the Set-MobileDeviceMailboxPolicy cmdlet. By default, access to both Windows file shares and Microsoft Windows SharePoint Services is allowed. You can block access to file shares and SharePoint by setting the -UNCAccessEnabled and -WSSAccessEnabled parameters to \$false.

If you specified that passwords are required, by default, simple passwords are not allowed. Additionally, by default, many device features are allowed. By using the TrueFalseParams shown in Listing 27-11, you can:

- Allow or disallow another device to share the device's Internet connection.
- Allow or disallow remote desktop connections.
- Allow or disallow the device to access email accounts other than Microsoft Exchange.
- Allow or disallow the device to access removable storage, such as memory cards.
- Allow or disallow the device to connect to a wireless network.
- Allow or disallow the device to connect to and synchronize with a desktop computer.
- Allow or disallow the device to connect to other devices using infrared.
- Allow or disallow the device to execute unsigned applications.
- Allow or disallow the device to install unsigned applications.
- Allow or disallow the device to use the built-in browser.
- Allow or disallow the device's built-in camera.

Use -MaxEmailBodyTruncationSize and -MaxEmailHTMLBodyTruncationSize to specify the maximum allowed size for email messages. Both parameter values are set in kilobytes. If a standard email message exceeds the MaxEmailBodyTruncationSize value, the message is truncated (clipped). If an HTML-formatted email message exceeds the MaxEmailHTMLBodyTruncationSize, the message is truncated (clipped).

If the policy allows devices to download attachments, the attachment has no default limit size. You can block attachment downloads by setting -AttachmentsEnabled to *\$false*. If you allow attachments and you want to limit the size of attachments that users can

download, you can specify the maximum allowed attachment size in kilobytes by using `-MaxAttachmentSize`.

For past calendar and email items, you can specify whether all calendar and mail items should be synced or only items from a specific period of time, such as the last two weeks. Use the `-MaxCalendarAgeFilter` and `-MaxEmailAgeFilter` parameters respectively. If you allow Bluetooth, you also can specify how the device can use Bluetooth. Set `-AllowBlueTooth` to `Allow` if you want to allow mobile devices to use Bluetooth in any mode or `HandsfreeOnly` to allow mobile devices to use Bluetooth only in hands-free mode. Set `-AllowBlueTooth` to `Disable` if you want to prevent devices from using BlueTooth.

LISTING 27-11 Set-MobileDeviceMailboxPolicy cmdlet syntax and usage

Syntax

```
Set-MobileDeviceMailboxPolicy -Identity MailboxPolicyId
[-AllowBluetooth <Disable | HandsfreeOnly | Allow>]
[-AllowSMIMEEncryptionAlgorithmNegotiation <BlockNegotiation |
OnlyStrongAlgorithmNegotiation | AllowAnyAlgorithmNegotiation>]
[-ApprovedApplicationList AppList] [-DevicePolicyRefreshInterval
Interval] [-DomainController FullyQualifiedName]
[-MaxAttachmentSize MaxSizeKB] [-MaxCalendarAgeFilter <All | TwoWeeks |
OneMonth | ThreeMonths | SixMonths>] [-MaxEmailAgeFilter <All |
OneDay | ThreeDays | OneWeek | TwoWeeks | OneMonth>]
[-MaxEmailBodyTruncationSize MaxSizeKB]
[-MaxEmailHTMLBodyTruncationSize MaxSizeKB] [-MaxInactivityTimeLock
InactiveTime] [-MaxPasswordFailedAttempts NumAttempts]
[-MinPasswordComplexCharacters MinComplexChars] [-MinPasswordLength
MinLength] [-MobileOTAUpdateMode <MajorVersionUpdates |
MinorVersionUpdates | BetaVersionUpdates>] [-Organization
OrgId] [-PasswordExpiration PasswordExp] [-PasswordHistory HistLength]
[-RequireEncryptionSMIMEAlgorithm <TripleDES | DES | RC2128bit |
RC264bit | RC240bit>] [-RequireSignedSMIMEAlgorithm <SHA1 | MD5>]
[-UnapprovedInROMApplicationList AppList] {OptionalTrueFalseParams}
```

{OptionalTrueFalseParams}

-AllowBrowser, -AllowCamera, -AllowConsumerEmail,
-AllowDesktopSync, -AllowExternalDeviceManagement, -AllowHTMLEmail,
-AllowInternetSharing, -AllowIrDA, -AllowMicrosoftPushNotifications,
-AllowMobileOTAUpdate, -AllowNonProvisionableDevices, -AllowPOPIMAPEmail,
-AllowRemoteDesktop, -AllowSimplePassword, -AllowSMIMESoftCerts,
-AllowStorageCard, -AllowTextMessaging, -AllowUnsignedApplications,
-AllowUnsignedInstallationPackages, -AllowWiFi,
-AlphanumericPasswordRequired, -AttachmentsEnabled,
-DeviceEncryptionEnabled, -IrmEnabled, -IsDefault, -PasswordEnabled,
-PasswordRecoveryEnabled, -RequireDeviceEncryption,
-RequireEncryptedSMIMEMessages, -RequireManualSyncWhenRoaming,
-RequireSignedSMIMEMessages, -RequireStorageCardEncryption,

-UNCAccessEnabled, -WSSAccessEnabled


Usage

Set-MobileDeviceMailboxPolicy -Identity "Device Policy for Executives"
-AllowNonProvisionableDevices \$false -AllowBluetooth HandsfreeOnly
-DeviceEncryptionEnabled \$true -PasswordRecoveryEnabled \$true
-RequireDeviceEncryption \$true -MaxAttachmentSize 5096
-MaxEmailBodyTruncationSize 10192 -MaxEmailHTMLBodyTruncationSize 10192

Assigning Mobile Device Mailbox Policies

The default mobile device mailbox policy is automatically applied by Exchange through implicit inheritance unless you assign a different non-default policy to a user. Any mailbox that has implicitly inherited policy automatically applies the currently-defined default policy and its settings. When you modify the default policy or configure a new default policy, you change the settings for all mailbox users that implicitly inherit the default policy.

To set a different policy as the default for new mailbox users, follow these steps:

1. In Exchange Admin Center, select **Mobile** in the Navigation menu, and then select **Mobile Device Mailbox Policies** to see a list of currently defined mobile device mailbox policies.
2. The current default policy has the value (default) as a suffix. To make another policy the default, select the policy you want to be the new default, and then select **Edit** ().
3. In the **Mobile Device Mailbox Policy** dialog box, select **This Is The Default Policy**, and then select **Save**.



Primary Device Mailbox Policy

general
security

*Name:
Primary Device Mailbox Policy

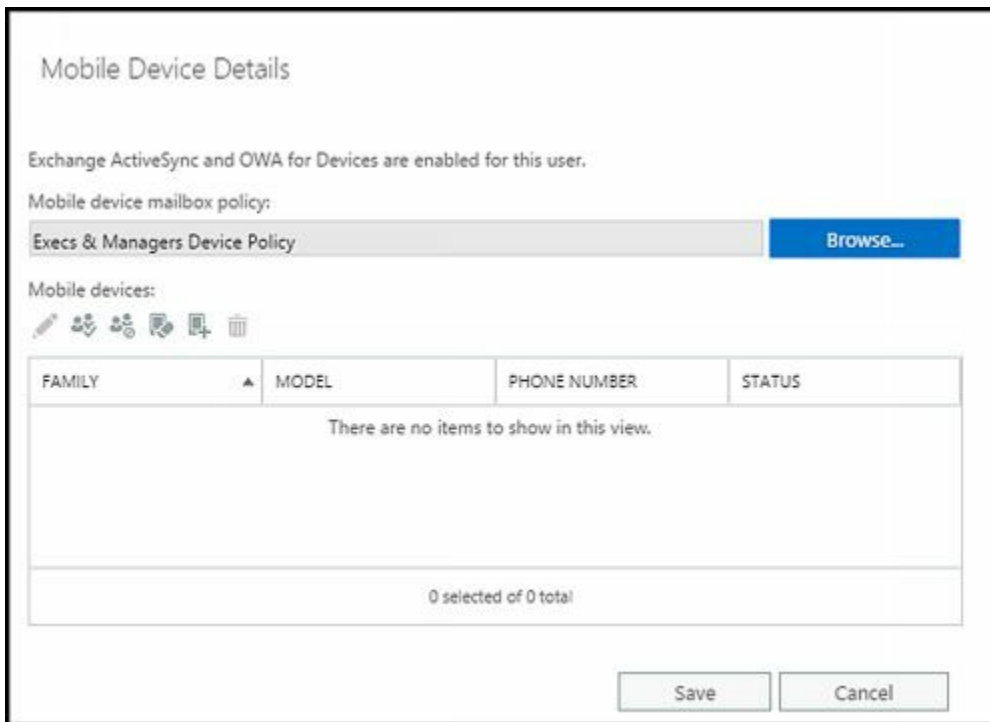
This is the default policy

Allow mobile devices that don't fully support these policies to synchronize

Save Cancel

To explicitly assign a policy to a mailbox, complete the following steps:

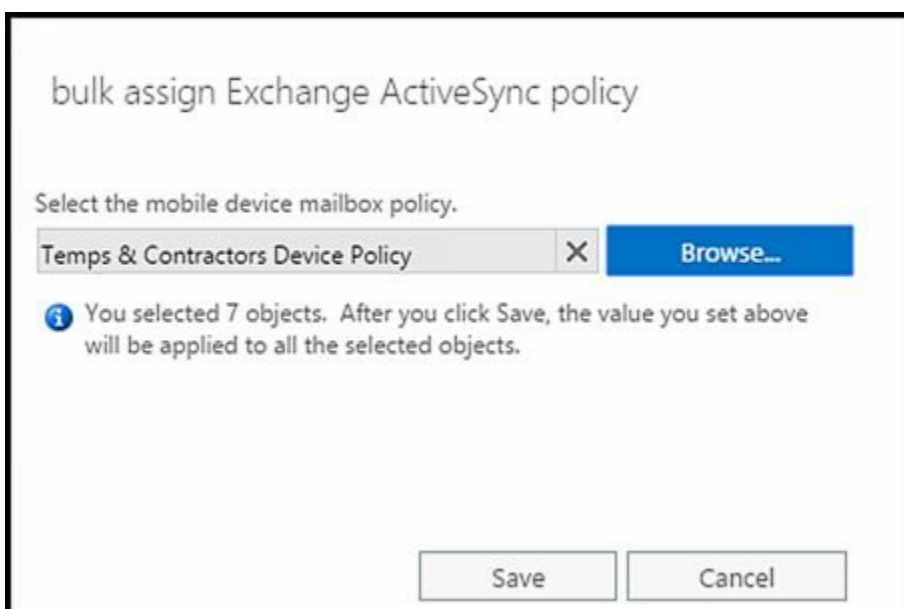
1. In Exchange Admin Center, select **Recipients** in the Navigation menu, and then select **Mailboxes**.
2. You should now see a list of users with Exchange mailboxes in the organization. Select the mailbox with which you want to work.
3. In the details pane, under **Mobile Devices**, select **View Details**.



4. In the Mobile Device Details dialog box, select Browse. Choose the policy to apply, and then select OK.
5. Click Save to apply your settings.

To explicitly assign a policy to multiple mailboxes, complete the following steps:

1. In Exchange Admin Center, select Recipients in the Navigation menu, and then select Mailboxes.
2. You should now see a list of users with Exchange mailboxes in the organization. Select multiple mailboxes by using the Shift or Ctrl keys.
3. In the details pane, scroll down. Under Exchange ActiveSync, select Update A Policy.
4. In the Bulk Assign.... dialog box, select Browse. Choose the policy to apply, and then select OK.
5. Click Save to apply your settings.



If you want mailbox users to use a mobile device mailbox policy other than the default, use the `-ActiveSyncMailboxPolicy` parameter of the `Set-CASMailbox` cmdlet to assign a policy directly to mailboxes. Listing 27-12 provides the syntax and usage.

LISTING 27-12 Assigning a Mobile Device Mailbox Policy to mailboxes

Syntax

```
Set-CASMailbox -Identity MailboxIdentity  
-ActiveSyncMailboxPolicy PolicyIdentity
```

Apply the policy to the mailbox user named MarkH

```
Set-CASMailbox -Identity "markh@tvpress.com"  
-ActiveSyncMailboxPolicy "Device Policy for Executives"
```

Apply the policy to every mailbox in the Exchange organization

```
Get-Mailbox -ResultSize Unlimited | Set-CASMailbox  
-ActiveSyncMailboxPolicy "Device Policy for Executives"
```

Apply the policy to every mailbox in the Sales database

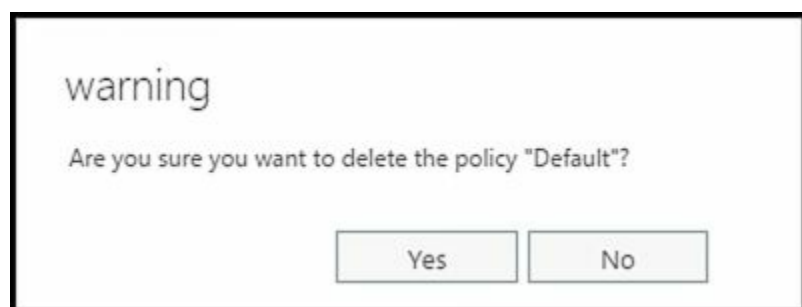
```
Get-MailboxDatabase "Sales" | Get-Mailbox -ResultSize Unlimited |  
Set-CASMailbox -ActiveSyncMailboxPolicy "Device Policy for Executives"
```

Apply the policy to all mailboxes in every mailbox database on MailboxServer18

```
Get-Mailbox -Server MailboxServer18 -ResultSize Unlimited |  
Set-CASMailbox -ActiveSyncMailboxPolicy "Device Policy for Executives"
```

Removing Mobile Device Mailbox Policies

When you no longer need a mobile device mailbox policy, you can remove it, provided that it isn't the current default policy. In the Exchange Admin Center, select the policy, and then select the Delete button. When prompted to confirm, click Yes to delete the policy. If users are assigned to the policy, they will stop using the policy and implicitly inherit the current default policy.



In Exchange Management Shell, you can remove a mobile device mailbox policy by using the `Remove-MobileMailboxPolicy` cmdlet. Listing 27-13 provides the syntax and usage.

LISTING 27-13 Remove-MobileMailboxPolicy cmdlet syntax and usage

Syntax

Remove-MobileMailboxPolicy -Identity Name [-DomainController DCName]
[-Force <\$true | \$false>]

Usage

Remove-MobileMailboxPolicy -Identity "Primary ActiveSync
Mailbox Policy"

Managing Device Access

To manage device access to Exchange, you can:

- [Block device access](#)
- [Define access rules](#)
- [Set access levels](#)
- [Set blocking thresholds](#)

Blocking Device Access

One way to prevent a device from synchronizing with Exchange is to put the device on the blocked mobile device list for the user's mailbox. The first step is to retrieve the ID of the device you want to prevent from syncing. Unfortunately, there's no way to retrieve the device ID before the user synchronizes the device with Exchange (unless you already know the device ID). If the user has synced the device already, you can get the device ID using:

```
Get-MobileDeviceStatistics -Mailbox ExchangeId  
-ActiveSync | fl DeviceID
```

Where *ExchangeId* is the email address or Exchange alias of the user, such as:

```
Get-MobileDeviceStatistics -Mailbox KaraH  
-ActiveSync | fl DeviceID
```

To prevent a device from synchronizing with Exchange, you must add the device to the `-ActiveSyncBlockedDeviceIDs` parameter list on the user's mailbox. To do this, run the following command:

```
Set-CASMailbox -Identity ExchangeID -ActiveSyncBlockedDeviceIDs  
@{Add="DeviceID"}
```

Where *ExchangeID* is the email address or Exchange alias for the mailbox user you want to prevent from using certain mobile devices, and *DeviceID* is the ID of the device to prevent from synchronizing with Exchange. If the device was previously on the user's allowed ActiveSync device list, you can remove the device from this list as well by using the following syntax:

```
Set-CASMailbox -Identity ExchangeID -ActiveSyncAllowedDeviceIDs  
@{Remove="DeviceID"}
```

NOTE As the blocked list has precedence over the allowed list, you technically don't have to remove the device from the allowed list. However, if someone accidentally resets the blocked list and you haven't removed the device from the allowed list, the user will be explicitly permitted to use the device to sync with Exchange.

Using Access Rules

Although you may sometimes want to manage device access for individual users, you'll probably prefer to define device access rules to control which device can and cannot sync with Exchange. To work with access rules, you'll use the following cmdlets:

- **Get-ActiveSyncDeviceAccessRule** Lists an access group of Exchange mobile devices along with their access level

```
Get-ActiveSyncDeviceAccessRule [-Identity AccessRuleId]
[-DomainController FullyQualifiedName] [-Organization OrgId]
```

- **Get-ActiveSyncDeviceClass** Lists mobile devices that have connected to Exchange by their type and model

```
Get-ActiveSyncDeviceClass [-Identity DeviceGroupId]
[-DomainController FullyQualifiedName] [-Filter FilterValues]
[-Organization OrgId] [-SortBy AttributeName]
```

- **New-ActiveSyncDeviceAccessRule** Defines an access group of Exchange mobile devices along with their access level

```
New-ActiveSyncDeviceAccessRule -AccessLevel <Allow | Block |
Quarantine> -Characteristic <DeviceType | DeviceModel | DeviceOS |
UserAgent> -QueryString Devices [-DomainController FullyQualifiedName]
[-Organization OrgId]
```

- **Remove-ActiveSyncDeviceAccessRule** Removes an existing device access rule

```
Remove-ActiveSyncDeviceAccessRule -Identity AccessRuleId
[-DomainController FullyQualifiedName]
```

- **Remove-ActiveSyncDeviceClass** Removes a device class from the list of mobile devices synchronizing with Exchange

```
Remove-ActiveSyncDeviceClass -Identity DeviceGroupId
[-DomainController FullyQualifiedName]
```

- **Set-ActiveSyncDeviceAccessRule** Sets the level of access for the ActiveSync Device Access rule

```
Set-ActiveSyncDeviceAccessRule -Identity AccessRuleId
[-AccessLevel <Allow | Block | Quarantine>]
[-DomainController FullyQualifiedName]
```

The following example creates access rules to block several different types of iOS 6.1 devices:

```
New-ActiveSyncDeviceAccessRule -querystring "iOS 6.1 10B142"
-characteristic DeviceOS -accesslevel block
```

```
New-ActiveSyncDeviceAccessRule -querystring "iOS 6.1 10B143"
-characteristic DeviceOS -accesslevel block
```

```
New-ActiveSyncDeviceAccessRule -querystring "iOS 6.1 10B144"
-characteristic DeviceOS -accesslevel block
```

Setting Access Levels and Blocking Thresholds

Another way to control device access is to define default access levels and blocking thresholds for ActiveSync devices. To do this, use the following cmdlets:

- **Get-ActiveSyncDeviceAutoblockThreshold** Lists the Autoblock settings for Exchange ActiveSync mobile devices

```
Get-ActiveSyncDeviceAutoblockThreshold [-Identity RuleName]  
[-DomainController FullyQualifiedName]
```

- **Set-ActiveSyncDeviceAutoblockThreshold** Modifies the autoblocking settings for mobile devices

```
Set-ActiveSyncDeviceAutoblockThreshold -Identity RuleName  
[-AdminEmailInsert MessageText] [-BehaviorTypeIncidenceDuration  
TimeSpan] [-BehaviorTypeIncidenceLimit Limit]  
[-DeviceBlockDuration TimeSpan] [-DomainController FullyQualifiedName]
```

- **Get-ActiveSyncOrganizationSettings** Lists the Exchange ActiveSync settings for the Exchange organization

```
Get-ActiveSyncOrganizationSettings [-Identity ExchangeOrgId]  
[-DomainController FullyQualifiedName] [-Organization OrgId]
```

- **Set-ActiveSyncOrganizationSettings** Modifies the Exchange ActiveSync settings for the Exchange organization

```
Set-ActiveSyncOrganizationSettings [-Identity ExchangeOrgId]  
[-AdminMailRecipients email1,email2,...emailN] [-DefaultAccessLevel  
<Allow | Block | Quarantine>] [-DomainController FullyQualifiedName]  
[-OtaNotificationMailInsert MessageText] [-UserMailInsert MessageText]
```


Chapter 28. Tracking and Logging Exchange Server 2016

As part of routine maintenance, you need to monitor Exchange Server to ensure that services and processes are functioning normally. Key components of any monitoring plan should include messaging tracking and protocol logging.

You use message tracking to monitor the flow of messages into, out of, and within an organization. With message tracking enabled, Exchange Server maintains daily log files, with a running history of all messages transferred within an organization. You use the logs to determine the status of a message, such as whether a message has been sent, has been received, or is waiting in the queue to be delivered. Because Exchange Server handles postings to public folders in much the same way as email messages, you can also use message tracking to monitor public folder usage.

Tracking logs can really save the day when you're trying to troubleshoot delivery and routing problems. The logs are also useful in fending off problem users who blame email for their woes. Generally speaking, users can't claim they didn't receive emails if you can find the messages in the logs. That said, if you use third-party applications that integrate with Outlook, those applications could potentially delete messages before the user sees them.

Protocol logging allows you to track Simple Mail Transfer Protocol (SMTP) communications that occur between servers as part of message routing and delivery. These communications could include both Exchange servers and non-Exchange servers. When non-Exchange servers send messages to an Exchange server, Exchange does the protocol logging of the communications.

You use protocol logging to troubleshoot problems with the Send and Receive connectors that are configured on Mailbox and Edge Transport servers. However, you shouldn't use protocol logging to monitor Exchange activity. This is primarily because protocol logging can be processor intensive and resource intensive, which means that an Exchange server may have to perform a lot of work to log protocol activity. The overhead required for protocol logging depends on the level of messaging activity on the Exchange server.

Configuring Message Tracking

By default, all Edge Transport and Mailbox servers perform message tracking. By setting the `-MessageTrackingLogEnabled` parameter of the `Set-TransportService` cmdlet to `$true` or `$false`, as appropriate, you can enable or disable message tracking on a per-server basis.

The following example disables message tracking on MailServer96:

```
Set-TransportService -Identity "MailServer96 "  
-MessageTrackingLogEnabled $false
```

TIP You can configure basic message tracking options in the Exchange Admin Center. To do this, select Servers in the Features pane, and then select the Servers tab. In the main pane, double-click the server you want to configure to display the related Properties dialog box. On the Transport Logs page select or clear the Enable Message Tracking Log check box. If you enable message tracking, you can enter the desired directory path for logging as well or accept the default setting.

Changing the Logging Location

Each Edge Transport and Mailbox server in your organization can have different message tracking settings that control the following:

- [Where logs are stored](#)
- [How logging is performed](#)
- [The maximum log size and maximum log directory size](#)
- [How long logs are retained](#)

By default, message tracking logs are stored in the `%ExchangeInstallPath%\TransportRoles\Logs\MessageTracking` directory. Generally, message tracking does not have high enough input/output activity to warrant a dedicated disk. However, in some high usage situations, you might want to move the tracking logs to a separate disk. Before you do this, however, you should create the directory you want to use and set the following required permissions:

- [Full Control for the server's local Administrators group](#)
- [Full Control for System](#)
- [Full Control for Network Service](#)

After you've created the directory and set the required permissions, you can change the location of the tracking logs by setting the `-MessageTrackingLogPath` parameter of the `Set-TransportService` cmdlet to the desired local directory. The following example sets the message tracking directory as `G:\Tracking` on MailServer96:

```
Set-TransportService -Identity "MailServer96"  
-MessageTrackingLogPath "G:\Tracking"
```

NOTE When you change the location of the message tracking directory, Exchange Server does not copy any existing tracking logs from the old directory to the new one. If you want all the logs to be in the same location, you should manually copy the old logs to the new location before you use Set-TransportService to change the message tracking directory.

Setting Logging Options

By default, all Edge Transport and Mailbox servers perform extended message tracking, which allows you to perform searches based on message subject lines, header information, sender, and recipient. If you don't want to collect information on potentially sensitive subject lines, you can disable subject line tracking by setting the – MessageTrackingLogSubjectLoggingEnabled parameter of the Set-TransportService cmdlet to \$false, as shown in the following example:

```
Set-TransportService -Identity "MailServer96"  
-MessageTrackingLogSubjectLoggingEnabled $false
```

Exchange Server continues to write to message tracking logs until a log grows to a specified maximum size, at which point Exchange Server creates a new log and then uses this log to track current messages. By default, the maximum log file size is 10 megabytes (MB). You can change this behavior by setting the – MessageTrackingLogMaxFileSize parameter to the desired maximum file size. You must qualify the desired file size by using B for bytes, KB for kilobytes, MB for megabytes, or GB for gigabytes. The following example sets the message log file size to 50 MB:

```
Set-TransportService -Identity "MailServer96"  
-MessageTrackingLogMaxFileSize "50MB"
```

Exchange Server overwrites the oldest message tracking logs automatically when tracking logs reach a maximum age or when the maximum log directory size is reached. By default, the maximum age is 30 days and the maximum log directory size is 1000 MB. You can use the –MessageTrackingLogMaxAge parameter to set the maximum allowed age in the following format:

DD.HH:MM:SS

where DD is the number of days, HH is the number of hours, MM is the number of minutes, and SS is the number of seconds. The following example sets the maximum age for logs to 90 days:

```
Set-TransportService -Identity "MailServer96"  
-MessageTrackingLogMaxAge "90.00:00:00"
```

You can set the maximum log directory size by using the – MessageTrackingLogMaxDirectorySize parameter. As with the maximum log file size, the qualifiers are B, KB, MB, and GB. The following example sets the maximum log directory size to 2 GB:

```
Set-TransportService -Identity "MailServer96"
```

-MessageTrackingLogMaxDirectorySize "2GB"

Searching the Tracking Logs

The tracking logs are useful in troubleshooting problems with routing and delivery. In Exchange Management Shell, you use `Get-MessageTrackingLog` to search through the message tracking logs. The related syntax is:

```
Get-MessageTrackingLog [-Start DateTime] [-Server ServerId]  
[-End DateTime] {AddtlParams}
```

```
{AddtlParams}  
[-DomainController DCName] [-EventId {"BadMail" | "Defer" | "Deliver" |  
"DSN" | "Expand" | "Fail" | "PoisonMessage" | "Receive" | "Redirect" |  
"Resolve" | "Send" | "Submit" | "Transfer"} ] [-InternalMessageId  
MessageTrackingLogId] [-MessageId MessageId] [-MessageSubject  
Subject] [-Recipients SMTPEmailAddress1, SMTPEmailAddress2,...]  
[-Reference ReferenceField] [-ResultSize NumEntriesToReturn]  
[-Sender SMTPEmailAddress]
```

These parameters allow you to search the message tracking logs in the following ways:

- By date
- By event ID
- By message ID
- By message subject
- By recipients
- By sender
- By server that processed the messages

Beginning an Automated Search

To begin a search, you must specify one or more of the previously listed identifiers as the search criteria. You must also identify a server in the organization that has processed the message in some way. This server can be the sender's server, the recipient's server, or a server that relayed the message.

You set the search criteria by using the following parameters:

- **-End** Sets the end date and time for the search.
- **-EventID** Specifies the ID of the event for which you want to search, such as a RECEIVE, SEND, or FAIL event.
- **-InternalMessageID** Specifies the ID of the message tracking log entries for which you want to search.
- **-MessageID** Specifies the ID of the message for which you want to search.
- **-MessageSubject** Specifies the subject of the message for which you want to search.
- **-Recipients** Sets recipient's SMTP email address or addresses to return
- **-Reference** Specifies the reference field value within the message for which you

want to search.

- **–Sender** Sets the sender's SMTP email address (listed in the From field of the message) to return.
- **–Server** Sets the name of the Transport or Mailbox server that contains the message tracking logs to be searched.
- **–Start** Sets the start date and time for the search.

Using the **–Start** and **–End** parameters, you can search for messages from a starting date and time to an ending date and time. Using the **–Server** parameter, you specify the server to search. Consider the following example:

```
Get-MessageTrackingLog -Start "05/25/2014 5:30AM" -End "05/30/2014 7:30PM"  
-Server MailServer96 -Sender tonyj@imaginedlands.com
```

In this example, you search for a messages sent by Tonyj@imaginedlands.com between 5:30 A.M. May 25, 2014 and 7:30 P.M. May 30, 2014.

IMPORTANT Keep in mind that only messages that match all of the search criteria you've specified are displayed. If you want to perform a broader search, specify a limited number of parameters. If you want to focus the search precisely, specify multiple parameters.

Reviewing Logs Manually

Exchange Server creates message tracking logs daily and stores them by default in the %ExchangeInstallPath%\TransportRoles\Logs\MessageTracking directory. For US-English, each log file is named by the date on which it was created, using one of the following formats:

- **MSGTRKYYYYMMDD-N.log**, such as **MSGTRK20140325-1.log** for the first log created on March 25, 2014 by the Transport service.
- **MSGTRKMAYYYYYMMDD-N.log**, such as **MSGTRKM20140325-1.log** for the first log created on March 25, 2014 and used with moderated messages for tracking approvals and rejections.
- **MSGTRKMDYYYYMMDD-N.log**, such as **MSGTRKM20140325-1.log** for messages delivered to mailboxes by the Mailbox Transport Delivery service.
- **MSGTRKMSYYYYMMDD-N.log**, such as **MSGTRKM20140325-1.log** for messages sent from mailboxes by the Mailbox Transport Submission service.

The message tracking logs store each message event on a single line. The information on a particular line is organized by comma-separated fields. Logs begin with a header that shows the following information:

- A statement that identifies the file as a message tracking log file
- The version of the Exchange Server that created the file
- The date on which the log file was created
- A comma-delimited list of fields contained in the body of the log file

Although not all of the fields are tracked for all message events, message event fields

and their meanings follow:

- Client-hostname** The hostname of the client making the request
- Client-ip** The IP address of the client making the request
- Connector-id** The identity of the connector used
- Custom-Data** Optional custom data that was logged
- Date-Time** The connection date and time
- Directionality** An indication of the source of the message
- Event-id** The type of event being logged, such as Submit
- Internal-message-id** An internal identifier used by Exchange to track a message
- Message-id** The message identifier
- Message-info** Any related additional information on the message
- Message-subject** The subject of the message
- Original-client-ip** The IP address for the original client
- Original-server-ip** The IP address for the original server
- Recipient-address** The email addresses of the message recipients
- Recipient-count** The total number of recipients
- Recipient-status** The status of the recipient email address
- Reference** The references, if any
- Related-recipient-address** The email addresses of any related recipients
- Return-path** The return path on the message
- Sender-address** The distinguished name of the sender's email address
- Server-hostname** The server on which the log entry was generated
- Server-ip** The IP address of the server on which the log entry was generated
- Source** The component for which the event is being logged, such as StoreDriver
- Source-context** The context of the event source
- Tenant-id** A tenant identifier
- Total-bytes** The total size of the message in bytes

You can view the message tracking log files with any standard text editor, such as Microsoft Notepad. You can also import the message tracking log files into a spreadsheet or a database. Follow these steps to import a message tracking log file into

Microsoft Office Excel:

1. With Excel 2013 or Excel 2016, select File and then select Open. On the Open panel, select Computer and then select Browse.
2. Use the Open dialog box to select the message tracking log file you want to open. Set the file type as All Files (*.*), select the log file, and then click Open.
3. The Text Import Wizard starts automatically. Click Next. On the Delimiters list, choose Comma. Click Next, and then click Finish.
4. The log file should now be imported. You can view, search, and print the message tracking log as you would any other spreadsheet.

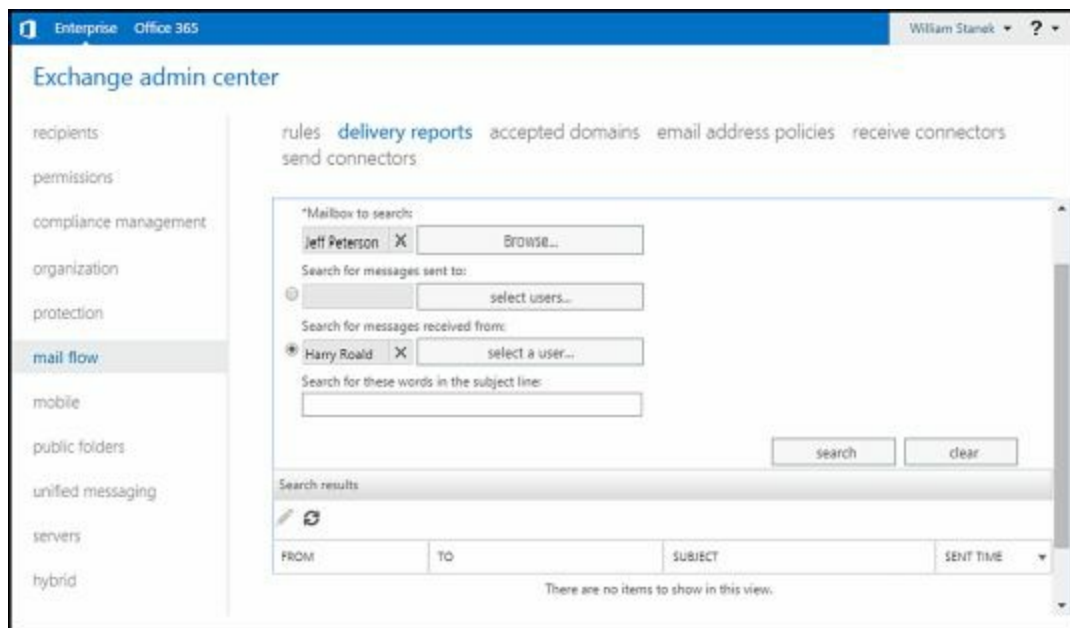
Searching the Delivery Status Reports

As part of message tracking, you can create delivery reports in Exchange Admin Center. Delivery reports allow you to search for the delivery status of messages sent to or from user's in your organization. In delivery reports, messages are listed by sender, recipients, and date and time sent. If subject line tracking is enabled, the subject line of messages is also included in reports.

You can track messages for up to 14 days after they were sent or received by completing the following steps:

1. In Exchange Admin Center, select Mail Flow in the Features pane, and then select Delivery Reports.

NOTE Only messages sent using SMTP, RPC or MAPI over HTTP or Outlook Web App can be tracked. Mail sent using POP3 or IMAP mail clients cannot be tracked.



2. Each delivery report is for messages sent to or from a specific mailbox. Under Mailbox To Search, click Browse. Select the mailbox to search, and then click OK.
3. Use the options provided to specify whether you want to search for messages sent from or to the mailbox you're searching. Keep the following in mind:
 - To find messages sent to specific users or groups from the mailbox you're searching, select Search For Messages Sent To, and then click Select Users. In the Select Users dialog box, select a user or group from the list, and then click Add. Repeat as necessary to add other users and groups. Click OK when you're finished.
 - To find all messages sent from the mailbox you're searching, select Search For Messages Sent To and then don't select any specify users or groups. By leaving the field blank, you create delivery reports for messages sent from the mailbox to anyone.
 - To find messages sent by a specific user to the mailbox you're searching, select

Search For Messages Received From, and then click Select A User. In the Select Members dialog box, select a user from the list, and then click Add. Click OK when you are finished. If you choose this option, you must select a user and cannot leave the field blank.

4. Optionally, if subject line tracking is enabled, you can restrict the search to messages with specific keywords in the subject line. In the Search For These Words... box, type one or more keywords to search for in the subject line of messages. To search for an exact phrase, enclose the phrase in quotation marks.
5. When you're ready to begin the search, click Search. If any matching messages are found, they are listed in the Search Results pane with the following fields:
 - **From** The display name, email address or alias of the person who sent the message.
 - **To** The display name, email address or alias of each message recipient.
 - **Sent** The date and time the message was sent.
 - **Subject** The subject line of the message.
6. View the delivery status and detailed delivery information for a message by selecting the message in the Search Results pane, and then selecting Details. When messages are sent to distribution groups, the details tell you the specific delivery status of each recipient in the group. When messages are moderated, the details tell you whether the moderator approved or rejected the message.

Configuring Protocol Logging

By default, protocol logging isn't enabled on custom connectors. As long as you know the identity of the custom connector with which you want to work, you can configure protocol logging for a specified connector. To retrieve a list of available Send and Receive connectors for a server, use the `Get-SendConnector` and `Get-ReceiveConnector` cmdlets, respectively. If you run either cmdlet without specifying additional parameters, a list of all available Send or Receive connectors is returned.

Enabling or Disabling Protocol Logging

You enable or disable protocol logging on a per-connector basis. For Send connectors, you use the `Set-SendConnector` cmdlet to enable protocol logging. For Receive connectors, you use the `Set-ReceiveConnector` cmdlet to enable protocol logging. Both cmdlets have a `-ProtocolLoggingLevel` parameter that you can set to `Verbose` to enable protocol logging or to `None` to disable protocol logging. Here is an example:

```
Set-ReceiveConnector -Identity "Corpsvr127\Custom Receive Connector"  
-ProtocolLoggingLevel 'Verbose'
```

Associated with the Transport service and the Front End Transport service on every Mailbox server is an implicitly created Send connector, referred to as the intra-organization Send connector. The Transport service on Mailbox servers uses the intra-organization Send connector to relay messages to other Transport servers in the Exchange organization. By default, the Transport service doesn't perform protocol logging on the intra-organization Send connector. You enable or disable protocol logging for the intra-organization Send connector on a Mailbox server by using the `-MailboxDeliveryConnectorProtocolLoggingLevel` parameter of the `Set-MailboxTransportService` cmdlet. Use `Verbose` to enable protocol logging or to `None` to disable protocol logging. Here is an example:

```
Set-MailboxTransportService -Identity MailServer96  
-MailboxDeliveryConnectorProtocolLoggingLevel 'Verbose'
```

The Frontend Transport service uses the intra-organization Send connector to relay messages to the Transport service on Mailbox servers. By default, the Front End Transport service performs protocol logging on the intra-organization Send connector. You enable or disable protocol logging for this Send connector by using the `-IntraOrgConnectorProtocolLoggingLevel` parameter of the `Set-FrontEndTransportService` cmdlet. Use `Verbose` to enable protocol logging or set to `None` to disable protocol logging. Here is an example:

```
Set-FrontEndTransportService -Identity MailServer26  
-IntraOrgConnectorProtocolLoggingLevel 'Verbose'
```

Setting Other Protocol Logging Options

Although you enable protocol logging on a per-connector basis, you configure the other protocol logging parameters on a per-server basis for either all Send connectors or all Receive connectors by using the Set-TransportService cmdlet. As it does with message tracking logs, Exchange Server overwrites the oldest protocol logs automatically when tracking logs reach a maximum age or when the maximum log directory size is reached. If you decide to move the protocol log directories, you should create the directories you want to use and then set the following required permissions:

- Full Control for the server's local Administrators group
- Full Control for System
- Full Control for Network Service

Because the parameters are similar to those for message tracking, I'll summarize the available parameters. The Send connector parameters for configuring protocol logging include:

- **SendProtocolLogMaxAge** Sets the maximum age for Send connector protocol logs. By default, set to 30.00:00:00.
- **SendProtocolLogMaxDirectorySize** Sets the maximum size for the Send connector protocol log directory. By default, set to 250 MB.
- **SendProtocolLogMaxFileSize** Sets the maximum size for Send connector protocol logs. By default, set to 10 MB.
- **SendProtocolLogPath** Sets the local file path for protocol logging of Send connectors. By default, set to %ExchangeInstallPath%\TransportRoles\Logs*ServerType* \ProtocolLogs\SmtpSend.

The Receive connector parameters for configuring protocol logging include:

- **ReceiveProtocolLogMaxAge** Sets the maximum age for Receive connector protocol logs. By default, set to 30.00:00:00.
- **ReceiveProtocolLogMaxDirectorySize** Sets the maximum size for the Receive connector protocol log directory. By default, set to 250 MB.
- **ReceiveProtocolLogMaxFileSize** Sets the maximum size for Receive connector protocol logs. By default, set to 10 MB.
- **ReceiveProtocolLogPath** Sets the local file path for protocol logging of Receive connectors. By default, set to %ExchangeInstallPath%\TransportRoles\Logs*ServerType* \ProtocolLogs\SmtpReceive.

In the default path for logs, the *ServerType* can be FrontEnd, Mailbox, or Hub. On Mailbox servers, you find these folders:

- Under FrontEnd\ProtocolLogs, you'll find logs for the Front End Transport service on Mailbox servers.
- Under Hub\ProtocolLogs, you'll find logs for the Transport service on Mailbox servers.
- Under Mailbox\ProtocolLogs, you'll find logs for the Mailbox Transport service on Mailbox servers.

TIP You can configure send and receive protocol log paths in the Exchange Admin Center. To do this, select Servers in the Features pane, and then select Servers. In the main pane, double-click the server you want to configure to display the related Properties dialog box. On the Transport Logs page, the Protocol log panel shows the current send and receive protocol log paths. You can specify the log file path by entering the desired directory path for logging or accept the default setting.

Managing Protocol Logging

When protocol logging is enabled, a Mailbox server or a transport server creates protocol logs daily. The protocol log stores each SMTP protocol event on a single line. The information on a particular line is organized by comma-separated fields. Logs begin with a header that shows the following information:

- A statement that identifies the file as either a Send connector protocol log or a Receive connector protocol log
- The date on which the log file was created
- The version of the Exchange Server that created the file
- A comma-delimited list of fields contained in the body of the log file

Although not all of the fields are tracked for all protocol events, SMTP event fields and their meanings include:

- **Connector-id** The distinguished name of the connector associated with the event.
- **Context** The context for the SMTP event.
- **Data** The data associated with the SMTP event.
- **Date-time** The date and time of the protocol event in a locale-specific format. For U.S. English, the format is *Y YYYY-MM-DDTHH:MM:SSZ*, such as 2014-06-05T23:30:59Z.
- **Event** The type of protocol event: + for Connect, – for Disconnect, > for Send, < for Receive, and * for Information.
- **Local-endpoint** The local endpoint of the SMTP session, identified by the Internet Protocol (IP) address and Transmission Control Protocol (TCP) port.
- **Remote-endpoint** The remote endpoint of the SMTP session, identified by the IP address and TCP port.
- **Sequence-number** The number of the event within an SMTP session. The first event has a sequence number of 0.
- **Session-id** The globally unique identifier of the SMTP session. Each event for a particular session has the same identifier.

You can view the protocol log files with any standard text editor, such as Notepad. You can also import the protocol log files into a spreadsheet or a database. Mailbox and transport servers store logs in either the

`%ExchangeInstallPath%\TransportRoles\Logs\ServerType\ProtocolLog\SmtpSend` or `%ExchangeInstallPath%\TransportRoles\Logs\ServerType\ProtocolLog\SmtpReceive` directory as appropriate for the type of server and connector being logged. For POP, IMAP, and other non-SMTP content aggregation, related logs are in the `%ExchangeInstallPath%\TransportRoles\Logs\ProtocolLog\HTTPClient` directory.

Each log file is named by the date on which it was created, using the format `SENDYYYYMMDD-N.log` or `RECVYYYYMMDD-N.log`, such as `SEND20160805-1.log` for the first Send connector log created on August 5, 2016. Additional protocol

logs are found in subdirectories of the %ExchangeInstallPath%\Logging directory. In the AddressBook Service subdirectory, you'll find logs for the Address Book service. In the RPC Client Access subdirectory, you'll find logs for Remote Procedure Calls for Client Access services.

Optimizing Protocol Logging for HTTP

@techjob

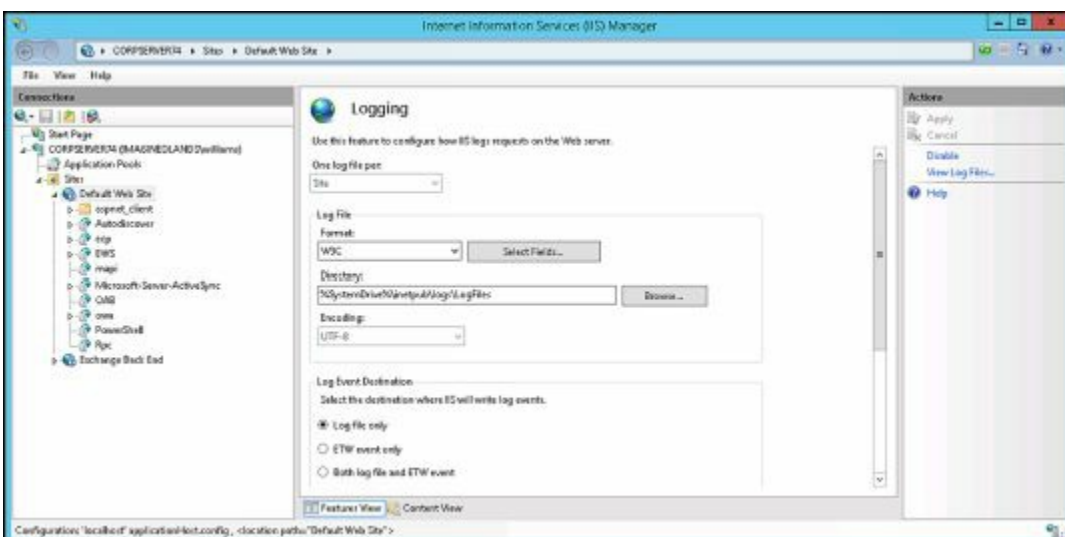
Mailbox servers have web-based applications and virtual directories that use Microsoft Internet Information Services (IIS) to provide the related services. In IIS, protocol logging for HTTP is a feature available when the HTTP Logging module is installed and logging is enabled. By default, this module is installed with IIS and enabled. The default configuration is to use one log file per website per day.

You can view and manage the logging settings by completing the following steps:

1. Start Internet Information Services (IIS) Manager. Start Server Manager, click Tools, and select Internet Information Services (IIS) Manager.

NOTE By default, IIS Manager connects to the services running on the local computer. If you want to connect to a different server, select the Start Page node in the left pane, and then click the Connect To A Server link. This starts the Connect To Server Wizard. Follow the prompts to connect to the remote server. Keep in mind that the Windows Remote Management Service must be configured and running on the remote server.

2. When you install the Mailbox role, the default website is created (or updated) to include the virtual directories and web-based applications used to provide front-end services for Exchange Server. If a server has the Mailbox role, a website named Exchange Back End is created and has virtual directories and web-based applications used to provide back-end services for Exchange Server. In IIS Manager, double-click the entry for the server with which you want to work, and then double-click Sites.
3. In the left pane, select the website that you want to manage, and then double-click Logging in the main pane to open the Logging feature.



4. If all logging options are dimmed and the server is configured for per-site logging, you can click **Enable** in the **Actions** pane to enable logging for this site. Otherwise, if logging is configured per server, you need to configure logging at the server level rather than at the site level; the procedure is similar.
 5. Use the **Format** selection list to choose one of the following log formats:
 - **W3C Extended Log File Format** Writes the log in ASCII text following the World Wide Web Consortium (W3C) extended log file format. Fields are space-delimited, and each entry is written on a new line. This style is the default. Using this option allows you to include extensive information about clients, servers, and connections.
 - **Microsoft IIS Log File Format** Writes the log in ASCII text following the IIS log file format. Fields are tab-delimited, and each entry is written on a new line. Using this option allows you to collect basic information about clients, servers, and connections.
 - **NCSA Common Log File Format** Writes the log in ASCII text following the National Center for Supercomputing Applications (NCSA) common log file format. Fields are space-delimited, and each entry is written on a new line. When you use this option, log entries are small because only basic information is recorded.
- TIP* W3C Extended Log File Format is the preferred logging format because you can record detailed information. Unless you're certain that another format meets your needs, you should use this format.
6. On the **Log File** panel, use the **Directory** text box to set the main folder for log files. By default, log files are written to a subdirectory of `%SystemDrive%\inetpub\logs\LogFiles`.
 7. On the **Log File Rollover** panel, select **Schedule** and then use the related selection list to choose a logging time period. In most cases, you'll want to create daily or weekly logs, so select either **Daily** or **Weekly**.
 8. If you selected **W3C**, click **Select Fields**, and then choose the fields that should be recorded in the logs. Click **Apply**.

Working with HTTP Protocol Logs

On Mailbox servers, HTTP protocol log files can help you detect and trace problems with HTTP, Outlook Web App, Microsoft Exchange ActiveSync, and Outlook Anywhere. By default, Exchange Server writes protocol log files to a subdirectory of `%SystemDrive%\inetpub\logs\LogFiles`. You can use the logs to determine the following:

- Whether a client was able to connect to a specified server and, if not, what problem occurred
- Whether a client was able to send or receive protocol commands and, if not, what error occurred
- Whether a client was able to send or receive data
- How long it took to establish a connection

- How long it took to send or receive protocol commands
- How long it took to send or receive data
- Whether server errors are occurring and, if so, what types of errors are occurring
- Whether server errors are related to Windows or to the protocol itself
- Whether a user is connecting to the server using the proper logon information

Most protocol log files are written as ASCII text. This means you can view them in Notepad or another text editor. You can import these protocol log files into Microsoft Office Excel in much the same way as you import tracking logs.

Log files, written as space-delimited or tab-delimited text, begin with a header that shows the following information:

- A statement that identifies the protocol or service used to create the file
- The protocol, service, or software version
- A date and timestamp
- A space-delimited or tab-delimited list of fields contained in the body of the log file

The name of the subdirectory used for logging depends on the number of websites hosted on a server. Typically, the W3SVC1 subdirectory is used for front-end logging and the W3SVC2 subdirectory is used for back-end logging.

Servers can have additional websites or may not have websites created in the expected order, such as when you deploy IIS prior to installing Exchange 2016. In this case, you'll want to confirm the identity of the logging subdirectory by using the following command:

```
Get-OwaVirtualDirectory -Server ServerID -ShowMailboxVirtualDirectories  
|fl identity, metabasepath
```

Where *ServerID* is the host name or fully-qualified domain name of the Exchange server to check, such as:

```
Get-OwaVirtualDirectory -Server MailServer21 -ShowMailboxVirtualDirectories  
|fl identity, metabasepath
```

The output will show the website identity and metabase path for the Outlook Web App (OWA) virtual directories created on the server. If the server has front-end and back-end virtual directories for OWA, the output will be similar to the following:

```
Identity    : MAILSERVER21\owa (Exchange Back End)  
MetabasePath : IIS://MAILSERVER21.imaginedlands.com/W3SVC/2/ROOT/owa
```

```
Identity    : MAILSERVER21\owa (Default Web Site)  
MetabasePath : IIS://MAILSERVER21.imaginedlands.com/W3SVC/1/ROOT/owa
```

In the output, note that the name of the associated website is shown in parenthesis as part of the identity and the subdirectory path can be extrapolated from the metabase path. Here, the back-end virtual directory is named Exchange Back End and has the associated subdirectory W3SVC2 (which is shown as W3SVC/2 in the metabase path). The front-end virtual directory is named Default Web Site and has the associated subdirectory W3SVC1 (which is shown as W3SVC/1 in the metabase path).

Using Connectivity Logging

Connectivity logging allows you to track the connection activity of outgoing message delivery queues. You use connectivity logging to troubleshoot problems with messages reaching their designated destination Mailbox server or recipient.

Configuring Connectivity Logging

By default, Exchange Server performs connectivity logging, creating connectivity logs when clients connect to the Front End Transport service on Mailbox servers and when clients are proxied or redirected to the Transport service on Mailbox servers where the related user's mailbox is stored. Exchange Server also creates connectivity logs for communications with the mailbox databases on a Mailbox server.

Generally, Exchange Server creates connectivity logs to track:

- When the Mailbox Transport Delivery receives SMTP messages from the Transport service and connects to local mailbox databases.
- When the Mailbox Transport Submission service connects to local mailbox databases to retrieve messages and submit them to the Transport service for delivery.

You manage connectivity logging for the Front End Transport service by using `Set-FrontEndTransportService`, the Transport service by using `Set-TransportService`, and the Mailbox Transport service by using `Set-MailboxTransportService`. With any of these cmdlets, you can enable or disable connectivity logging for the service by setting the `–ConnectivityLogEnabled` parameter to `$true` or `$false`, as appropriate. The following example disables connectivity logging for the Transport service on `MailServer96`:

```
Set-TransportService -Identity "MailServer96"  
-ConnectivityLogEnabled $false
```

TIP You can use the Exchange Admin Center to configure basic logging options for the Transport service (but not for other services). To do this, select Servers in the Features pane, and then select Servers. In the main pane, double-click the server you want to configure to display the related Properties dialog box. On the Transport Logs page select or clear the Enable Connectivity Logging check box. If you enable connectivity logging, you can specify the log file path, and then click OK.

The Front End Transport service, the Transport service, and the Mailbox Transport service can have different connectivity logging settings:

- Use the `–ConnectivityLogMaxAge` parameter to set the maximum log file age. The default maximum age is `30.00:00:00`.
- Use the `–ConnectivityLogMaxDirectorySize` parameter to set the maximum log directory size. The default maximum log directory size is 250 MB.
- Use the `–ConnectivityLogMaxFileSize` parameter to set the maximum log file size. The

default maximum log file size is 10 MB.

- Use the `–ConnectivityLogPath` parameter to move the log directory to a new location. The default logging directory depends on the service with which you are working.

As it does with other logs, Exchange Server overwrites the oldest connectivity logs automatically when tracking logs reach a maximum age or when the maximum log directory size is reached. If you decide to move the protocol log directories, you should create the directories you want to use and set the following required permissions:

- Full Control for the server's local Administrators group
- Full Control for System
- Full Control for Network Service

Working with Connectivity Logs

Exchange Server creates connectivity logs daily and stores them in the `%ExchangeInstallPath%\TransportRoles\Logs\ServerType\Connectivity` directory. Each log file is named by the date on which it was created, using the format `CONNECTLOGYYYYMMDD-N.log`, such as `CONNECTLOG20160319-1.log` for the first connectivity log created on March 19, 2016.

The connectivity log stores outgoing queue connection events on a single line. The information on a particular line is organized by comma-separated fields. Logs begin with a header that shows the following information:

- A statement that identifies the file as a connectivity log
- The date on which the log file was created
- The version of Exchange Server that created the file
- A comma-delimited list of fields contained in the body of the log file

Although not all of the fields are tracked for all outgoing queue connection events, connectivity logging fields and their meanings are:

- **Date-time** The date and time of the outgoing queue connection event.
- **Session** The globally unique identifier of the SMTP session. Each event for a particular session has the same identifier. For Messaging Application Programming Interface (MAPI) sessions, this field is blank.
- **Destination** The name of the destination Mailbox server, smart host, or domain.
- **Direction** The direction of the event: + for Connect, – for Disconnect, > for Send, and < for Receive.
- **Description** The data associated with the event, including the number and size of messages transmitted, Domain Name Server (DNS) name resolution information, connection success messages, and connection failure messages.

You can view the connectivity log files with any standard text editor, such as Notepad. You can also import the connectivity log files into a spreadsheet or a database, as discussed previously.

In the default path for logs, the `ServerType` can be `FrontEnd`, `Mailbox`, or `Hub`. On

Mailbox servers, you find these folders:

- Under FrontEnd\Connectivity, you'll find logs for the Front End Transport service.
- Under Hub\Connectivity, you'll find logs for the Transport service.
- Under Mailbox\Connectivity, you'll find a Submission subdirectory containing logs for the Mailbox Transport Submission service, and a Delivery subdirectory containing logs for the Mailbox Transport Delivery service.

Chapter 29. Maintaining Exchange Server 2016

You must maintain Microsoft Exchange Server 2016 to ensure proper mail flow and recoverability of message data. Routine maintenance should include monitoring event logs, services, servers, and resource usage as well as regular monitoring of Exchange queues. Mailbox and Edge Transport servers use queues to hold messages while they are processing them for routing and delivery. If messages remain in a queue for an extended period, problems could occur. For example, if an Exchange server is unable to connect to the network, you'll find that messages aren't being cleared out of queues. Because you can't be on-site 24 hours a day, you may want to set alerts to notify you when problems occur.

Monitoring Events, Services, Servers, and Resource Usage

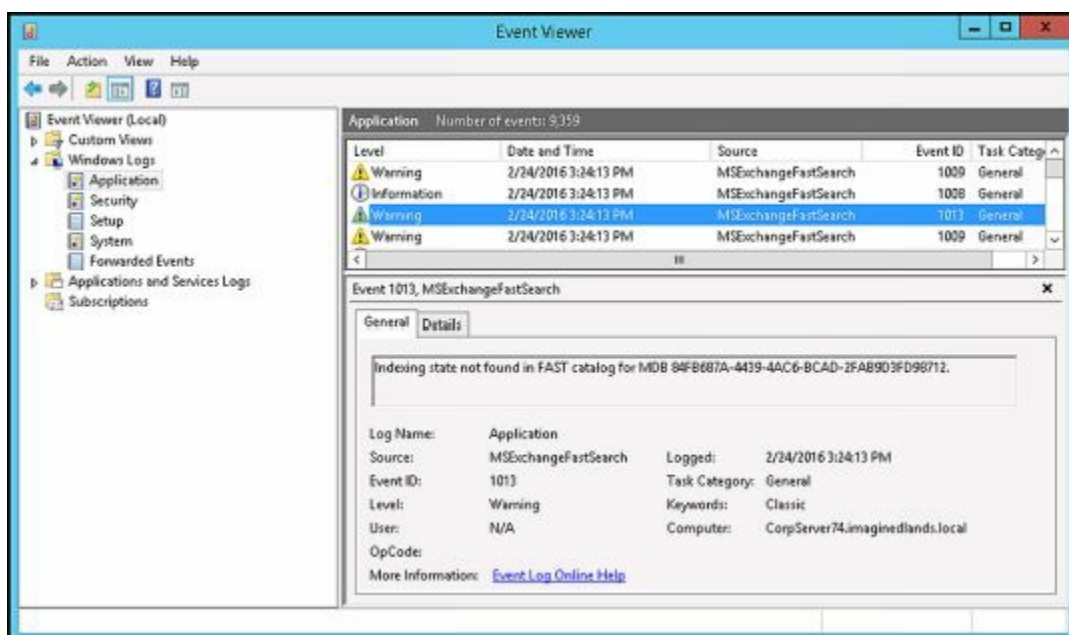
As discussed in Chapter 30, "Troubleshooting Exchange Server 2016," Exchange 2016 includes a built-in monitoring and problem resolution architecture that can resolve many types of issues automatically. The automated responders will take recovery actions automatically, which can include restarting services and restarting servers. However, the automated processes won't detect all issues, which is why you must routinely monitor Exchange as well.

Viewing Events

System and application events generated by Exchange Server are recorded in the Windows event logs. The primary log that you'll want to check is the application log. In this log, you'll find the key events recorded by Exchange Server services. Keep in mind that related events might be recorded in other logs, including the directory service, DNS server, security, and system logs. For example, if the server is having problems with a network card and this card is causing message delivery failures, you'll have to use the system log to pinpoint the problem.

You access the application log by completing the following steps:

1. In Server Manager, click Tools, and then select Event Viewer.
2. If you want to view the logs on another computer, in the console tree, right-click the Event Viewer entry, and choose Connect To Another Computer from the shortcut menu. You can now choose the server for which you want to manage logs.
3. Double-click the Windows Logs node. You should now see a list of logs.
4. Select the Application log.



Entries in the main panel of Event Viewer provide an overview of when, where, and how an event occurred. To obtain detailed information on an event, select its entry. The event level precedes the date and time of the event. Event levels include the following:

- **Information** An informational event, generally related to a successful action.
- **Warning** Details for warnings are often useful in preventing future system problems.
- **Error** An error such as the failure of a service to start.
- **Critical** A critical error such as the failure of an essential clustering component.

In addition to level, date, and time, the summary and detailed event entries provide the following information:

- **Source** The application, service, or component that logged the event.
- **Event ID** An identifier for the specific event.
- **Task Category** The category of the event, which is sometimes used to further describe the related action.
- **User** The user account that was logged on when the event occurred.
- **Computer** The name of the computer on which the event occurred.
- **Description** In the detailed entries, this event entry provides a text description of the event.
- **Data** In the detailed entries, this event entry provides any data or error code output created by the event.

Use the event entries to detect and diagnose Exchange performance problems. Exchange-related event sources include the following:

- **ESE** Helps you track activities related to the Extensible Storage Engine (ESE) used by the Information Store. Watch for logging and recovery errors, which might indicate a problem with a database or a recovery action. For example, Event ID 300 indicates the database engine initiated recovery steps; Event ID 301 indicates the database engine has begun replaying a log file for a mailbox database; and Event ID 302 indicates the database engine has successfully completed recovery steps. If you want to track the status of online defragmentation, look for Event ID 703. Additional related sources include ESENT and ESE Backup.
- **MSEExchange Antimalware, MSEExchange Antispam, MSEExchange Anti-spam Update** Helps you track activities related to anti-malware and anti-spam agents. When you've configured Microsoft Exchange to use Microsoft Update to retrieve anti-spam updates, watch for errors regarding update failure. You might need to change the Microsoft Update configuration or the way updates are retrieved.
- **MSEExchange Assistants, MSEExchangeMailboxAssistants** Helps you track activities related to the Microsoft Exchange Mailbox Assistants service. The Microsoft Exchange Mailbox Assistants service performs background processing and maintenance of mailboxes. Watch for processing errors, which can indicate database structure problems.
- **MSEExchange EdgeSync** Helps you track activities related to the Edge

Synchronization processes. The Microsoft Exchange EdgeSync service uses the Exchange Active Directory Provider to obtain information about the Active Directory topology. If the service cannot locate a suitable domain controller, the service fails to initialize and edge synchronization fails as well.

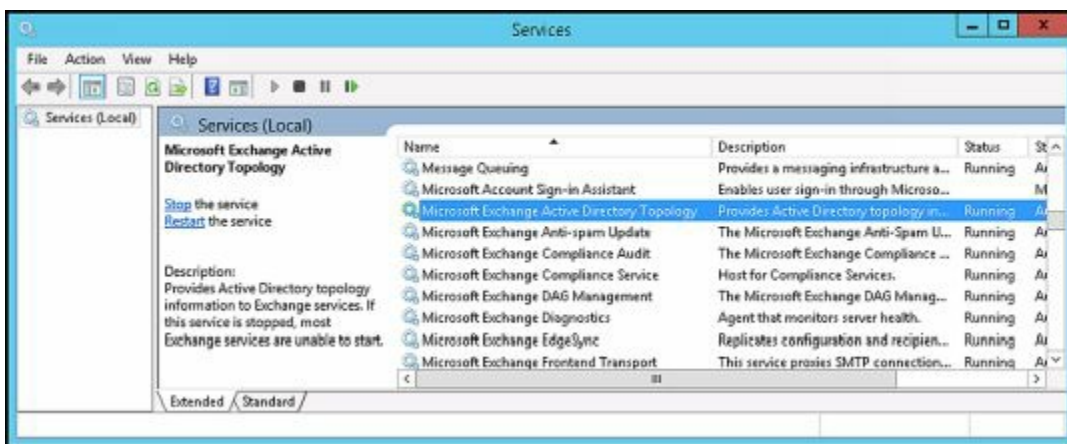
- **MSEExchange TransportService, MSEExchange Transport** Helps you track activities related to the Microsoft Exchange Transport service and message transport in general. Watch for errors that can indicate issues with storage or shadow redundancy. Related sources include MSEExchangeDelivery and MSEExchangeTransportDelivery for tracking the Mailbox Transport Delivery service, and MSEExchangeSubmission and MSEExchangeTransportSubmission for tracking the Mailbox Transport Submission service.
- **MSEExchangeADAccess** Helps you track activities related to the Exchange Active Directory Provider, which is used for retrieving information for Active Directory and performing the DNS lookups that Exchange uses to locate domain controllers and global catalog servers. Watch for topology discovery failures and DNS lookup failures, which can indicate problems with the DNS configuration as well as with the Active Directory site configuration.
- **MSEExchangeDiagnostics, MSEExchangeHM** Helps you track activities related to the Microsoft Exchange Diagnostics service and the Microsoft Exchange Health Manager, respectively. With diagnostics, watch for errors related to low disk space and low available memory. With the health manager, watch for errors related to the working processes. Also MSEExchangeHMHost.
- **MSEExchangeFrontEndTransport, MSEExchange Front End HTTP Proxy** Help you track activities related to Front End Transport service and Front End HTTP proxying of web applications, respectively. Related sources include MSEExchangeOWA for tracking the Outlook Web App, MSEExchange Web Services for tracking Exchange Web Services, and MSEExchange RPC Over HTTP Autoconfig for tracking the configuration of Outlook Anywhere.
- **MSEExchangeIS** Helps you track activities related to the Microsoft Exchange Information Store service and mailbox databases. If a user is having problems logging on to Exchange, you might see multiple logon errors. You might also see a lot of logon errors if someone is trying to hack into an Exchange mailbox. Watch also for errors related to high availability.
- **MSEExchangeRepl** Helps you track activities related to Active Manager and database failover. Watch for errors related to mounting, moving, or unmounting databases.

Managing Essential Services

Most of Exchange Server's key components run as system services. If an essential service stops, its related functionality will not be available and Exchange Server won't work as expected. When you are troubleshooting Exchange Server problems, you'll want to check to ensure that essential services are running as expected early in your troubleshooting process. To manage system services, you can use the Services console or the Services node in the Computer Management console. You can start and work with

the Services console by completing the following steps:

1. In Server Manager, click Tools, and then select Services.
2. If you want to manage the services on another computer, right-click the Services entry in the console tree, and select Connect To Another Computer on the shortcut menu. You can now choose the system with which you want to work.
3. You'll now see the available services. Services are listed by:
 - **Name** The name of the service.
 - **Description** A short description of the service and its purpose.
 - **Status** The status of the service. If the entry is blank, the service is stopped.
 - **Startup Type** The startup setting for the service.
 - **Log On As** The account the service logs on as. The default in most cases is the local system account.



TIP Any service that has a startup type of Automatic should have a status of Started. If a service has a startup type of Automatic and the status is blank, the service is not running and you should start it (unless another administrator has stopped it to perform maintenance or troubleshooting).

If a service is stopped and it should be started, you need to restart it. If you suspect a problem with a service, you can perform try to diagnose the problem as discussed in Chapter 30, “Troubleshooting Exchange Server 2016,” and might also want to stop and then restart it. To start, stop, or restart a service, complete the following steps:

1. Access the Services console.
2. Right-click the service you want to manage, and then select Start, Stop, or Restart, as appropriate.

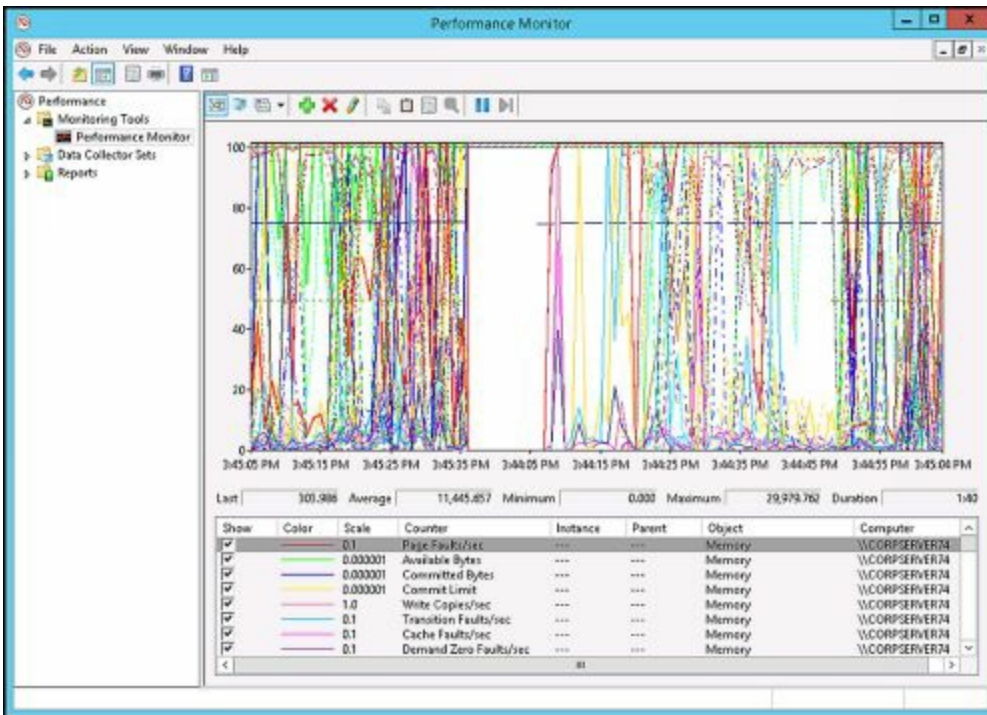
After you start or restart a service, you should check the event logs to see if there are errors related to the service. Any related errors you find might help you identify why the service wasn't running.

Keep in mind that Exchange 2016 automatically restarts services that are found to not be responding or otherwise need restarting as part of the Managed Availability architecture. The automated processes can also reset IIS and restart servers. Although these automated processes work well, they won't always resolve service issues as

quickly as you could by manually intervening.

Monitoring Messaging Components

When you are troubleshooting or optimizing a server for performance, you can use performance monitoring to track the activities of Exchange messaging components. Performance Monitor graphically displays statistics for the set of performance parameters you've selected for display. These performance parameters are referred to as *counters*. Performance Monitor displays information only for the counters you're tracking. Thousands of counters are available, and these counters are organized into groupings called *performance objects*.



When you install Exchange Server 2016 on a computer, Performance Monitor is updated with a set of objects and counters for tracking Exchange performance. These objects and counters are registered during setup in the Win32 performance subsystem and the Windows registry. You'll find several hundred related performance objects for everything from the Microsoft Exchange Active Manager to the Microsoft Exchange Journaling Agent to Microsoft Exchange Outlook Web App.

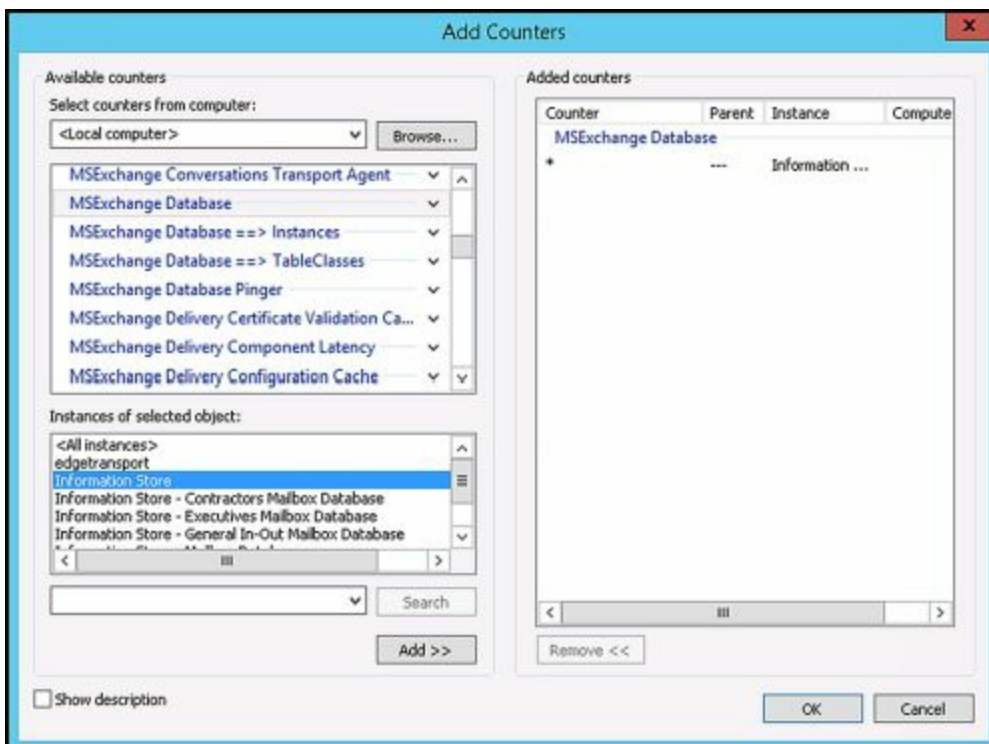
You can select which counters you want to monitor by completing the following steps:

1. In Server Manager, select Tools, and then select Performance Monitor. Next, select the Performance Monitor entry in the left pane.
2. The Performance Monitor tool has several views and view types. Ensure that you are viewing current activity by clicking View Current Activity on the toolbar or pressing Ctrl+T. You can switch between the view types (Line, Histogram Bar, and Report) by clicking the Change Graph Type button or pressing Ctrl+G.
3. To add counters, click Add on the toolbar or press Ctrl+N. This displays the Add Counters dialog box.
4. In the Select Counters From Computer combo box, enter the Universal Naming

Convention (UNC) name of the Exchange server with which you want to work, such as \\MailServer96 , or leave it at the default setting of <Local computer> to work with the local computer.

NOTE You need to be at least a member of the Performance Monitor Users group in the domain or the local computer to perform remote monitoring. When you use performance logging, you need to be at least a member of the Performance Log Users group in the domain or the local computer to work with performance logs on remote computers.

5. In the Available Counters panel, performance objects are listed alphabetically. If you select an object entry by clicking it, all related counters are selected. If you expand an object entry, you can see all the related counters and you can then select individual counters by clicking them. For example, you can expand the entry for the MExchangeTransport Database object and then select the DataRow clones/sec, Database Connections Current and MailItem new/sec counters.



6. When you select an object or any of its counters, you see the related instances, if any. Choose All Instances to select all counter instances for monitoring separately. Choose _total to view a single combined value reflecting data for all available instances. Or select one or more counter instances to monitor. For example, when you select MExchangeIS Store, you'll find separate instances for each database on the server and you could select an individual database to specifically track that database.
7. When you've selected an object or a group of counters for an object as well as the object instances, click Add to add the counters to the graph. Repeat steps 5 through 6 to add other performance parameters.
8. Click OK when you're finished adding counters. You can delete counters later by clicking their entry in the lower portion of the Performance window, and then

clicking Delete.

Using Performance Alerting

Data Collector Sets are used to collect performance data. When you configure Data Collector Sets to alert you when specific criteria are met, you are using performance alerting. Windows performance alerting provides a fully automated method for monitoring server performance and reporting when certain performance thresholds are reached. You can use performance alerting to track the following:

- Memory usage
- CPU utilization
- Disk usage
- Messaging components

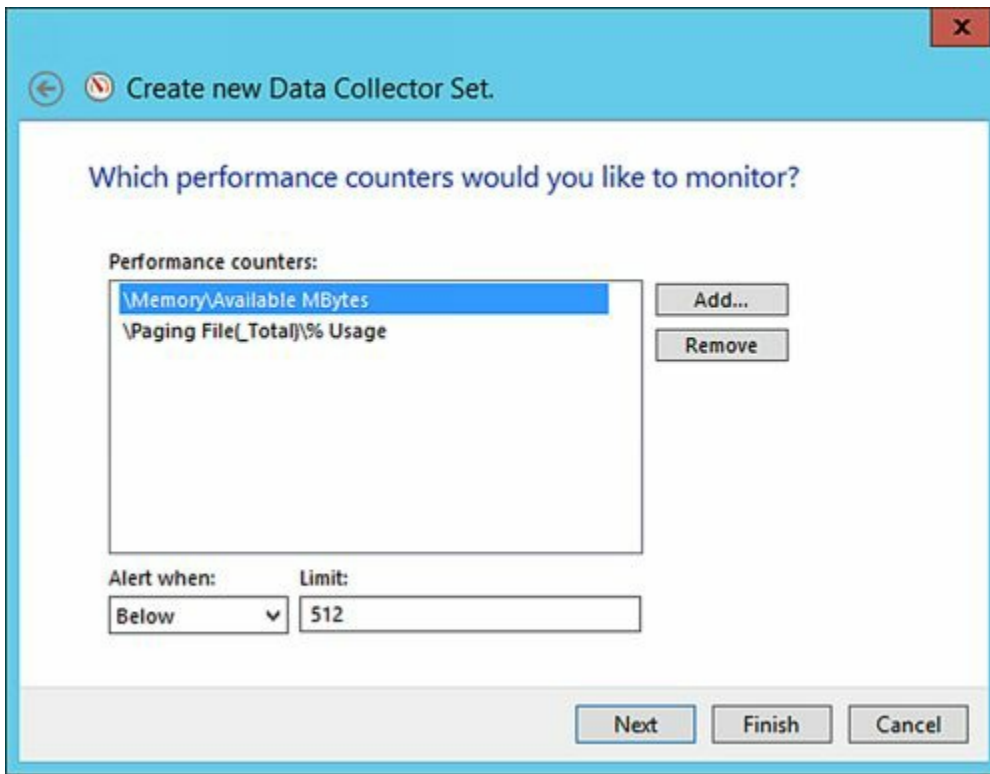
Using notifications, you can then provide automatic notification when a server exceeds a threshold value.

Tracking Memory Usage

Physical and virtual memory is critical to normal system operation. When a server runs low on memory, system performance can suffer and message processing can grind to a halt. To counter this problem, you should configure performance alerting to watch memory usage. You could then increase the amount of virtual memory available on the server or add more random access memory (RAM) as needed. However, keep in mind that increasing virtual memory isn't something you should do without careful planning.

You configure a memory alert by completing the following steps:

1. In Server Manager, select Tools, and then select Performance Monitor. Next, expand the Data Collector Sets node, and then select User Defined. You should see a list of current alerts (if any) in the right pane.
2. Right-click the User-Defined node in the left pane, point to New, and then choose Data Collector Set.
3. In the Create New Data Collector Set wizard, type a name for the Data Collector Set, such as **Memory Usage Alert**. Select the Create Manually option, and then click Next.
4. On the What Type Of Data Do You Want To Include page, the Create Data Logs option is selected by default. Select the Performance Counter Alert option, and then click Next.
5. On the Which Performance Counters Would You Like To Monitor page, click Add. This displays the Add Counter dialog box. Because you are configuring memory alerts, expand the Memory object in the Performance Object list. Select Available MBytes by clicking it, and then click Add.



6. Expand the Paging File object in the Performance Object list. Click %Usage. In the Instances Of Selected Object panel, select _Total, and then click Add. Click OK.
7. On the Which Performance Counters Would You Like To Monitor page, you'll see the counters you've added. In the Performance Counters panel, select Available MBytes, set the Alert When list to Below, and then enter a Limit value that is approximately 5 to 8 percent of the total physical memory (RAM) on the server for which you are configuring alerting. For example, if the server has 8 GB of RAM, you could set the value to 512 MB to alert you when the server is running low on available memory.
8. In the Performance Counters panel, select %Usage. Set the Alert When list to Above, and then type **98** as the Limit value. This ensures that you are alerted when more than 98 percent of the paging file is being used.
9. Click Next, and then click Finish. This saves the Data Collector Set and closes the wizard.
10. In the left pane, under User Defined, select the related Data Collector Set, and then double-click the data collector for the alert in the main pane. This displays the data collector Properties dialog box.
11. On the Alerts tab, use the Sample Interval options to set a sample interval. The sample interval specifies when new data is collected. Don't sample too frequently, however, because you'll use system resources and might cause the server to seem unresponsive. By default, Performance Monitor checks the values of the configured counters every 15 seconds. A better value might be once every 10 to 30 minutes. Generally, you'll want to track performance periodically over several hours at a minimum and during a variety of usage conditions.
12. If you want to log an event rather than be alerted every time an alert limit is

reached, on the Alert Action tab, select the Log An Entry In The Application Event Log check box. Selecting this option ensures that an event is logged when the alert occurs but does not alert you via the console. Click OK to close the Properties dialog box.

By default, alerting is configured to start manually. To start alerting, select the User Defined node in the left pane, click the alert in the main pane to select it, and then click the Start button on the toolbar. If you later want to stop getting alerts, right-click the alert in the main pane, and then select Stop.

Tracking CPU Utilization

You can use a CPU utilization alert to track the usage of a server's CPUs. When CPU utilization is too high, Exchange Server can't effectively process messages or manage other critical functions. As a result, performance can suffer greatly. For example, CPU utilization at 100 percent for an extended period of time can be an indicator of serious problems on a server. To recover, you might need to use Task Manager to end the process or processes with high CPU utilization, or you might need to take other corrective actions to resolve the problem, such as closing applications you are running while logged on to the server.

You'll also want to closely track process threads that are waiting to execute. A relatively high number of waiting threads can be an indicator that a server's processors need to be upgraded.

You configure a CPU utilization alert by completing the following steps:

1. In Server Manager, select Tools, and then select Performance Monitor. Next, expand the Data Collector Sets node, and then select User Defined. You should see a list of current alerts (if any) in the right pane.
2. Right-click the User-Defined node in the left pane, point to New, and then choose Data Collector Set.
3. In the Create New Data Collector Set wizard, type a name for the Data Collector Set, such as **CPU Utilization Alert**. Select the Create Manually option, and then click Next.
4. On the What Type Of Data Do You Want To Include page, the Create Data Logs option is selected by default. Select the Performance Counter Alert option, and then click Next.
5. On the Which Performance Counters Would You Like To Monitor page, click Add. This displays the Add Counter dialog box. Because you are configuring CPU alerts, expand the Processor object in the Performance Object list. Click % Processor Time. In the Instances Of Selected Object panel, select _Total, and then click Add.
6. Expand the System object in the Performance Object list. Click Processor Queue Length, and then click Add. Click OK.
7. On the Which Performance Counters Would You Like To Monitor page, you'll see

the counters you've added. Select % Processor Time. Then set the Alert When list to Above, and type **98** as the Limit value. This ensures that you are alerted when processor utilization is more than 98 percent.

8. In the Performance Counters panel, select Processor Queue Length. Then set the Alert When list to Above, and type **3** as the Limit value. This ensures that you are alerted when more than three processes are waiting to execute, which can be an indicator that a server's processors need to be upgraded.
9. Click Next, and then click Finish. This saves the Data Collector Set and closes the wizard.
10. Finish configuring the alert by following steps 10 through 12 under "Tracking Memory Usage" earlier in this chapter.

Tracking Disk Usage

Exchange Server uses disk space for data storage, logging, tracking, and virtual memory. To ensure ample disk space is always available, Exchange Server monitors free disk space. If free disk space drops below specific thresholds, Exchange will gracefully shut itself down. When Exchange is in this state, it is likely that data could get lost. To prevent serious problems, you should monitor free disk space closely on all drives used by Exchange Server.

You'll also want to track closely the number of system requests that are waiting for disk access. A relatively high value for a particular disk can affect server performance and is also a good indicator that a disk is being overutilized or that there may be some problem with the disk. To resolve this problem, you'll want to try to shift part of the disk's workload to other disks, such as by moving databases, logs, or both.

You configure disk usage alerting by completing the following steps:

1. In Server Manager, select Tools, and then select Performance Monitor. Next, expand the Data Collector Sets node, and then select User Defined. You should see a list of current alerts (if any) in the right pane.
2. Right-click the User-Defined node in the left pane, point to New, and then choose Data Collector Set.
3. In the Create New Data Collector Set wizard, type a name for the Data Collector Set, such as **Disk Usage Alert**. Select the Create Manually option and then click Next.
4. On the What Type Of Data Do You Want To Include page, the Create Data Logs option is selected by default. Select the Performance Counter Alert option, and then click Next.
5. On the Which Performance Counters Would You Like To Monitor page, click Add. This displays the Add Counter dialog box. Because you are configuring disk alerts, expand the LogicalDisk object in the Performance Object list. Click % Free Space. In the Instances Of Selected Object panel, select all individual logical disk instances that you want to track. Do not select `_Total` or `<All`

Instances>. Click Add.

6. Expand the PhysicalDisk object in the Performance Object list. Click Current Disk Queue Length. In the Instances Of Selected Object panel, select all individual physical disk instances except _Total, and then click Add. Click OK.
7. On the Which Performance Counters Would You Like To Monitor page, you'll see the counters you've added. Select the first logical disk instance, set the Alert When list to Below, and then type **5** as the Limit value. This ensures that you are alerted when available free space is less than 5 percent. Repeat this procedure for each logical disk.
8. In the Performance Counters panel, select the first physical disk instance, set the Alert When list to Above, and then type **2** as the Limit value. This ensures that you are alerted when more than two system requests are waiting for disk access. Repeat this procedure for each physical disk.
9. Click Next, and then click Finish. This saves the Data Collector Set and closes the wizard.
10. Finish configuring the alert by following steps 10 through 12 under "Tracking Memory Usage" earlier in the chapter.

Working with Queues

As an Exchange administrator, it's your responsibility to monitor Exchange queues regularly. Mailbox and Edge Transport servers use queues to hold messages while they are processing them for routing and delivery. If messages remain in a queue for an extended period, problems could occur. For example, if an Exchange server is unable to connect to the network, you'll find that messages aren't being cleared out of queues.

Understanding Exchange Queues

Queues are temporary holding locations for messages that are waiting to be processed, and Exchange Server 2016 uses an Extensible Storage Engine (ESE) database for queue storage. Exchange Server 2016 uses the following types of queues:

- **Submission queue** The submission queue is a persistent queue that is used by the Exchange Categorizer (a transport component) to temporarily store all messages that have to be resolved, routed, and processed by transport agents. All messages that are received by a transport server enter processing in the submission queue. Messages are submitted through SMTP-receive, the Pickup directory, or the store driver. Each transport server has only one submission queue. Messages that are in the submission queue cannot be in other standard queues at the same time.

Edge Transport servers use the Categorizer to route messages to the appropriate destinations. Mailbox servers use the Categorizer to expand distribution lists, to identify alternative recipients, and to apply forwarding addresses. After the Categorizer retrieves the necessary information about recipients, it uses that information to apply policies, route the message, and perform content conversion. After categorization, the transport server moves the message to a delivery queue or to the Unreachable queue.

- **Mailbox delivery queue** Mailbox delivery queues hold messages that are being delivered to a Mailbox server by using encrypted Exchange RPC. Only Mailbox servers have mailbox delivery queues, and they use the queue to temporarily store messages that are being delivered to mailbox recipients whose mailbox data is stored on a Mailbox server that is located in the same site as the Mailbox server. Mailbox servers have one mailbox delivery queue for each destination Mailbox server associated with messages currently being routed. After queuing the message, the Mailbox server delivers the messages to the distinguished name of the mailbox database.
- **Relay queue** Relay queues hold messages that are being relayed to another server. Only Mailbox servers have relay queues, and they use the queue to temporarily store messages that are being delivered to mailbox recipients whose mailbox data is being relayed through a connector, designated expansion server, or non-SMTP gateway. Mailbox servers have one relay queue for each connector, designated expansion server, or non-SMTP gateway. After queuing a message, the Mailbox server relays the

message.

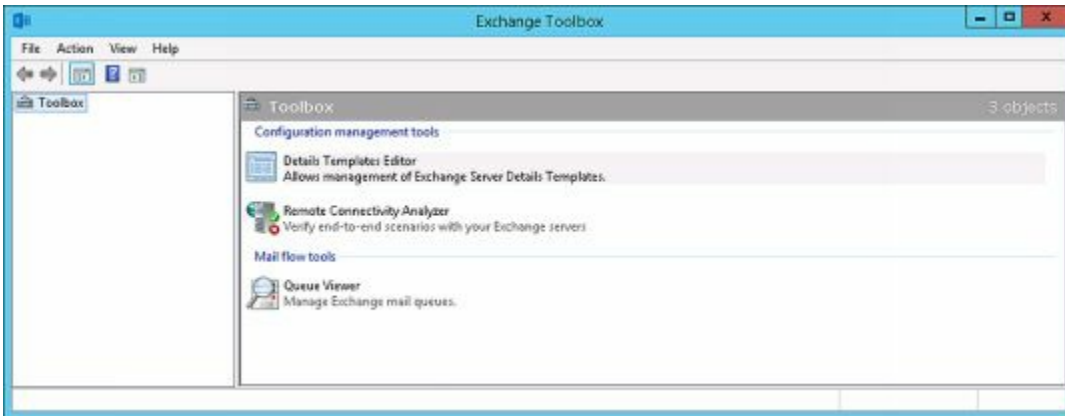
- **Remote delivery queue** Remote delivery queues hold messages that are being delivered to a remote server by using SMTP. Edge Transport servers can have remote delivery queues, and they use the queue to temporarily store messages that are being routed to remote destinations. On an Edge Transport server, these destinations are external SMTP domains or SMTP connectors. Edge Transport servers have one remote delivery queue for each remote destination associated with messages currently being routed. After queuing the message, the transport server delivers it to the appropriate server, smart host, IP address, or Active Directory site. Mailbox servers running Exchange 2016 do not have remote delivery queues.
- **Poison message queue** The poison message queue is used to hold messages that are detected to be potentially harmful to Exchange Server 2016 after a server failure. Messages that contain errors that are potentially fatal to Exchange Server 2016 are delivered to the poison message queue. Each Mailbox server has one poison message queue, as does each Edge Transport server. Although this queue is persistent, it typically is empty and, as a result, is not displayed in queue viewing interfaces. By default, all messages in the poison message queue are suspended and can be manually deleted.
- **Shadow redundancy queue** The shadow redundancy queue is used to prevent the loss of messages that are in transit by storing queued messages until the next transport server along the route reports a successful delivery of the message. If the next transport server doesn't report successful delivery, the message is resubmitted for delivery. This queue is nonpersistent. Mailbox and Edge Transport servers have one for each hop to which the server delivered the primary message.
- **Safety Net queue** The Safety Net queue keeps a redundant copy of messages that have been successfully processed by a Mailbox server. If a message needs to be redelivered, a Mailbox server can resend the message from the Safety Net queue. Each Mailbox server has one primary Safety Net queue and one shadow Safety Net queue. These queues are nonpersistent.
- **Transport dumpster queue** The transport dumpster queue is used to hold messages that are being delivered. This queue is nonpersistent. Edge Transport servers have one queue for each Active Directory site. Mailbox servers do not have a transport dumpster queue.
- **Unreachable queue** The unreachable queue contains messages that cannot be routed to their destinations. Each Mailbox server has one unreachable queue, as does each Edge Transport server. Although this queue is persistent, it typically is empty and, as a result, is not displayed in queue viewing interfaces.

When a transport server receives a message, a transport mail item is created and saved in the appropriate queue within the queue database. Exchange Server assigns each mail item a unique identifier when it stores the mail item in the database. If a mail item is being routed to more than one recipient, the mail item can have more than one destination and, in this case, there is a routed mail item for each destination. A routed mail item is a reference to the transport mail item, and it is the routed mail item that

Exchange queues for delivery.

Accessing the Queue Viewer

Using Queue Viewer, you can track message queues and mail flow. On any computer in which you've installed the Exchange management tools, you'll be able to access the Queue Viewer from the Exchange Toolbox. Open Exchange Toolbox from Start and then double-click Queue Viewer.

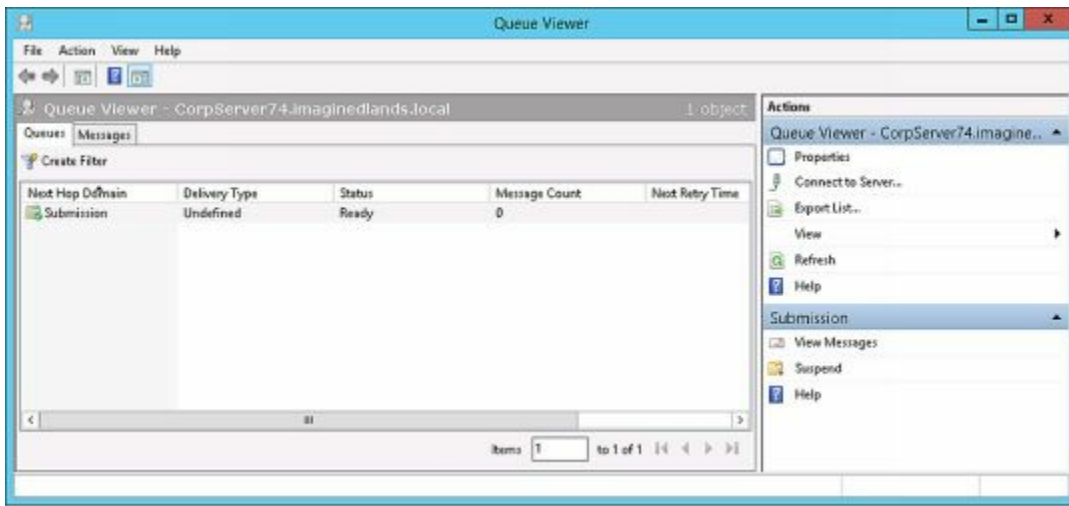


By default, the Queue Viewer connects to the queuing database on the local server (if applicable). To connect to a different server, on the Actions pane, select Connect To Server. In the Connect To Server dialog box, click Browse. Select the Exchange Server with which you want to work, and then click OK. Finally, click Connect.



The Queue Viewer provides an overview of the status of each active queue, including the following information:

- A folder icon indicates an active state.
- A folder icon with a green check mark indicates the queue has a ready status.
- A folder icon with a blue button and a small down arrow indicates a retry state.
- A folder icon with a red exclamation point indicates a warning state, such as Not Available or Error.



Managing Queues

You usually won't see messages in queues because they're processed and routed quickly. Messages come into a queue, Exchange Server performs a lookup or establishes a connection, and then Exchange Server either moves the message to a new queue or delivers it to its destination.

Understanding Queue Summaries and Queue States

Messages remain in a queue when there's a problem or if they have been suspended by an administrator. To check for problem messages, use the Queue Viewer to examine the number of messages in the queues. If you see a queue with a consistent or growing number of messages, there might be a problem. Again, normally, messages should come into a queue and then be processed quickly. Because of this, the number of messages in a queue should gradually decrease over time as the messages are processed, provided no new messages come into the queue.

Whenever you click the Queues tab in the Queue Viewer, you get a summary of the currently available queues for the selected server. Although queue summaries provide important details for troubleshooting message flow problems, you do have to know what to look for. The connection status is the key information to look at first. This value tells you the state of the queue. States you'll see include the following:

- **Active** An active queue has messages that are being transported.
- **Ready** A ready queue is needed to allow messages to be transported. When queues are ready, they can have a connection allocated to them.
- **Retry** A connection attempt has failed and the server is waiting to retry.
- **Suspended** The queue is suspended, and none of its messages can be processed for routing. Messages can enter the queue, but only if the Exchange Categorizer is running. You must resume the queue to resume normal queue operations.

Administrators can choose to enable or disable connections to a queue by right-clicking the queue and selecting Suspend. If a queue is suspended, it's unable to route and deliver messages.

You can change the queue state to Ready by right-clicking the queue and selecting Resume. When you do this, Exchange Server should immediately enable the queue, which allows messages to be routed and delivered. If a queue is in the retry state, you can force an immediate retry by using the Retry command.

Other summary information that you might find useful in troubleshooting include the following:

- **Delivery Type** Tells you what type of recipient messages are being queued for delivery.
- **Next Hop Domain** Tells you the next destination of a delivery queue. For mailbox

delivery, relay, and remote delivery queues, this field tells you the next hop domain. Messages queued for delivery to an EdgeSync server list the associated site and destination, such as EdgeSync–Default-First-Site To Internet.

- **Message Count** Tells you the total number of messages waiting in the queue. If you see a large number, you might have a connectivity or routing problem.
- **Next Retry Time** When the connection state is Retry, this column tells you when another connection attempt will be made. You can click the Retry command to attempt a connection immediately.
- **Last Retry Time** When the connection state is Retry, this column tells you when the last retry attempt was made.
- **Last Error** Tells you the error code and details of the last error to occur in a particular queue. This information can help you determine why a queue is having delivery problems.

You can add or remove columns by using the Add/Remove Columns dialog box. Display this dialog box by choosing View in the Actions pane and then selecting Add/Remove Columns.

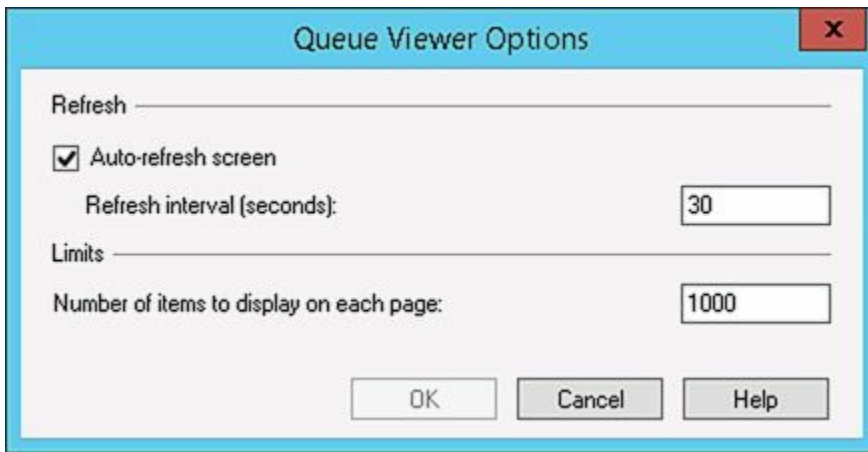
REAL WORLD Queue Viewer uses Windows PowerShell to perform all actions, including displaying and refreshing queue data. To display the commands Queue Viewer is using, choose View in the Actions pane, and then select View Exchange Management Shell Command Log.

Refreshing The Queue View

Use the queue summaries and queue state information to help you find queuing problems, as discussed in the “Understanding Queue Summaries and Queue States” section earlier in this chapter. By default, the queue view is refreshed every 30 seconds, and the maximum number of message items that can be listed on each page is 1,000.

To change the viewing options, follow these steps:

1. In the Queue Viewer, on the View menu, click Options.
2. To turn off automatic refresh, clear the Auto-Refresh Screen check box. Otherwise, enable automatic refresh by selecting the Auto-Refresh Screen check box.
3. In the Refresh Interval text box, type a specific refresh rate in seconds.
4. Type the desired maximum number of messaging items to be displayed per page in the Number Of Items To Display text box. Click OK.

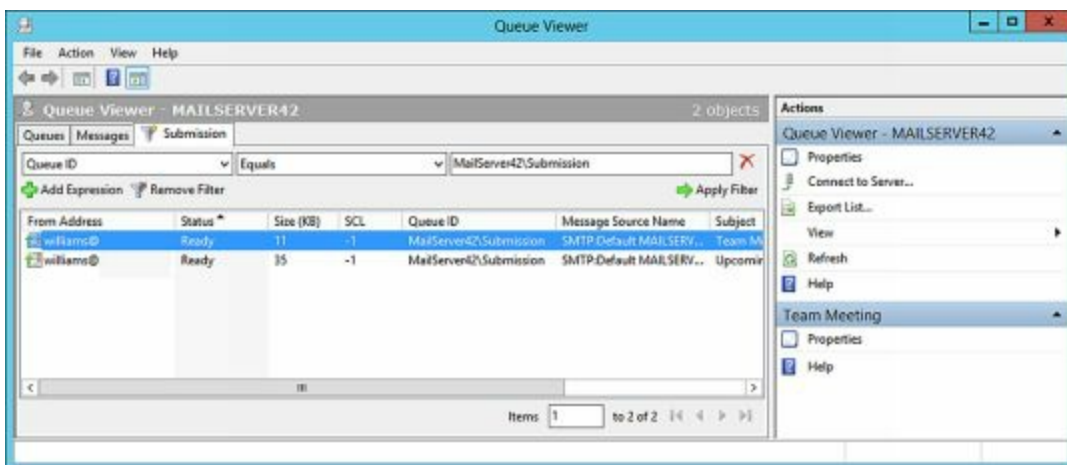


Working with Messages In Queues

To manage queues, you must enumerate messages. This process allows you to examine queue contents and perform management tasks on messages within a particular queue.

The easiest way to enumerate messages is to do so in sets of 1,000. To display the first 1,000 messages in a queue, follow these steps:

1. On the Queues tab in the Queue Viewer, you should see a list of available queues. Double-click a queue to enumerate the first 1,000 messages.
2. After you enumerate messages in a queue, you can examine message details by double-clicking the entries for individual messages. This enumerates the first 1,000 messages in the selected queue by filtering the message queues based on the queue identifier of the selected queue.



You can also create a filter to search for specific types of messages. To do this, follow these steps:

1. Double-click the queue with which you want to work. This enumerates the first 1,000 messages in the selected queue by filtering the message queues based on the queue identifier of the selected queue.
2. Click Add Expression. Use the first selection list to specify the field you want to use for filtering messages. You can filter messages by the following criteria: Date Received, Expiration Time, From Address, Internet Message ID, Last Error, Message Source Name, Queue ID, SCL, Size (KB), Source IP, Status, and

Subject.

3. Use the second selection list to specify the filter criteria. The available filter criteria depend on the filter field and include Equals, Does Not Equal, Contains, Does Not Contain, Greater Than, and Less Than.
4. Use the text box provided to specify the exact criteria to match. For example, if you are filtering messages using the Status field, you might want to see all messages in which the Status field equals Retry.
5. To apply the new filter criteria, click Apply Filter.

Forcing Connections to Queues

In many cases, you can change the queue state to Ready by forcing a connection. Simply right-click the queue, and then select Retry. When you do this, Exchange Server should immediately enable connections to the queue, and this should allow messages to be routed to and delivered from the queue.

Suspending and Resuming Queues

When you suspend a queue, all new message transfer activity out of that queue stops and only messages being processed will be delivered. This means that messages can continue to enter the queue, but no new messages will leave it. To restore normal operations, you must resume the queue.

You suspend and resume a queue by completing the following steps:

1. On the Queues tab in the Queue Viewer, you should see a list of available queues. Right-click a queue, and then select Suspend.
2. When you're done troubleshooting, right-click the queue, and then select Resume.



Another way to suspend messages in a queue is to do so selectively. In this way, you can control the transport of a single message or several messages that might be causing problems on the server. For example, if a large message is delaying the delivery of other messages, you can suspend that message until other messages have left the queue. Afterward, you can resume the message to resume normal delivery.

To suspend and then resume individual messages, complete the following steps:

1. On the Messages tab in the Queue Viewer, you should see a list of queued messages.
2. Right-click the message you want to suspend, and then select Suspend. You can select multiple messages by using Shift and Ctrl.
3. When you're ready to resume delivery of the message, right-click the suspended message, and then select Resume.

Deleting Messages from Queues

You can remove messages from queues if necessary. To do this, follow these steps:

1. On the Messages tab in the Queue Viewer, you should see a list of queued messages.
2. Right-click the message you want to remove. You can select multiple messages by using Shift and Ctrl, and then right-click. Select one of the following options from the shortcut menu:
 - **Remove (With NDR)** Deletes the selected messages from the queue, and notifies the sender with a nondelivery report (NDR)
 - **Remove (Without Sending NDR)** Deletes the message or messages from the queue without sending an NDR to the sender
3. When prompted, click Yes to confirm the deletion.

Deleting messages from a queue removes them from the messaging system permanently. You can't recover the deleted messages.

Chapter 30. Troubleshooting Exchange Server 2016

Microsoft Exchange Server 2016 is critically important to your organization and to be a successful Exchange administration you need to know how to diagnose and resolve problems as quickly as possible. Throughout this book, I've discussed techniques you can use to configure, maintain, and troubleshoot Microsoft Exchange Server 2016. In this chapter, I discuss additional techniques you can use to perform comprehensive troubleshooting.

Troubleshooting Essentials

Mailbox servers running Exchange 2016 can experience many types of issues that require troubleshooting to resolve. These issues can range from performance problems, to denied logins, to service outages. To help you resolve problems as they occur, you need a solid understanding of Exchange architecture, which I've covered throughout this book as part of the core discussion. Now let's look at architecture components specific to maintaining, diagnosing, and resolving Exchange services.

Tracking Server Health

In Exchange Server 2016, the Managed Availability architecture is used to automatically detect and correct many types of system problems with a goal of helping to ensure the overall availability of Exchange services. Managed Availability is implemented as part of the Mailbox server role. All servers running Exchange 2016 have this architecture.

As part of Managed Availability, hundreds of probes, monitors, and responders are running constantly on Exchange 2016 to analyze, monitor, and maintain services. If a problem is identified, it often can be fixed automatically. Managed Availability has three asynchronous components:

- **Probe Engine** Takes measurements on the server and collects data samples. The collected data flows to the monitor engine.
- **Monitor Engine** Uses the measurements and collected data to determine the status of Exchange services and components. The processed data flows to the responder engine.
- **Responder Engine** Takes recovery actions based on unhealthy states reported by the monitor engine. If automated recovery is unsuccessful, escalates by issuing event log notifications.

By delving deeper into the Managed Availability architecture, you can get a better understanding of how the automated monitoring and response processes work. The workflow has three phases:

- **Sampling** The probe engine checks the state of Exchange services and components according to specific probes. Each probe has a top-level identifier and one or more related probe definitions. Each probe definition identifies the name of the associated probe, the health set to which the probe belongs, the target resource being tracked, a recurrence interval, and a timeout value.
- **Detection** The monitor engine analyzes the sampled data and issues alerts related to changes in the state of Exchange services and components according to specific monitors. Each monitor has a top-level identifier and one or more related monitor definitions. Each monitor definition identifies the name of the associated monitor, the health set to which the monitor belongs, and a sample mask that specifies the top level identifier for related probes.

- **Recovery** The responder engine responds to unhealthy states identified in alerts. Each responder has an associated responder definition that identifies the recovery action to be taken, the name of the responder, the target resource that will be acted on, and an alert mask that specifies the top-level identifier for related monitors.

NOTE Rather than list each associated monitor or probe, Managed Availability components use name masking. Here, a top-level identifier is provided and then used as a mask to identify the related monitors and probes.

Collections of monitors are grouped together in health sets. Exchange 2016 has health sets for everything from Microsoft ActiveSync to User Throttling. Each health set has a number of associated monitors. As part of automated recovery, responders use the alerts issued by monitors to take recovery actions. There are three levels of recovery:

- **Tier 1** Provides the initial recovery response. As an initial response to an unhealthy state, responders typically will try to restart the service that uses the affected components.
- **Tier 2** Provides more advanced and customized recovery response. If restarting the service doesn't resolve the issue, the monitor state is escalated to the next level. The action or actions taken at this level to recover depend on the component but could include failover, bug checking, re-initialization of components to bring them back online, and more.
- **Tier 3** Uses the escalate responder to issue event log notifications regarding the problem. If you've installed the Exchange Server 2016 Management Pack, escalated issues are sent to Microsoft System Center Operations Manager via the event logs as well.

Although designed to resolve many typical problems, Managed Availability cannot resolve every problem and this escalation is built into the architecture. As part of diagnosing and resolving problems, you can check the status of monitors and health sets by using:

- **Get-HealthReport** Details the state and health of Exchange resources, monitors, and services.

```
Get-HealthReport -Identity ServerID [-GroupSize SizeOfRollup]  
[-HaImpactingOnly <$true | $false>] [-HealthSet HealthSet]  
[-MinimumOnlinePercent MinToDegraded>]  
[-RollupGroup <$true | $false>]
```

- **Get-ServerHealth** Returns the state of monitored resources as well as alert values.

```
Get-ServerHealth -Identity ServerID [-HaImpactingOnly <$true | $false>]  
[-HealthSet HealthSet]
```

To check the state of resources, enter the following command:

```
Get-ServerHealth -Identity ServerID
```

Where ServerID is the host name or fully-qualified name of the Exchange server to check, such as:

```
Get-ServerHealth -Identity MailServer42
```

In the following sample, I've omitted the server name and server component columns from the default output:

State	Name	TargetResource	HealthSetName	AlertValue
Online	AutodiscoverProxy...	MSEExchangeAutoDis...	Autodiscover...	Healthy
Online	ActiveSyncProxyTe...	MSEExchangeSyncApp...	ActiveSync.P...	Healthy
Repairing	ECPPProxyTestMonitor	MSEExchangeECPAppPool	ECP.Proxy	Unhealthy

REALWORLD Often when you work with Exchange Management Shell, you'll find that the output is too long for the default screen buffer size or that the output has too many columns for the default window size. Because of this, I prefer to use a screen buffer height of 2999 and width of 120, along with a window width of 120 and height of 74. This makes Exchange Management Shell easier to work with. If you are using Windows 8.1 or later or Windows Server 2012 or later, you'll find that you can't customize all of these settings from the Start screen. Instead, right-click the tile for the shell on the Start screen, and then select Open File Location. This opens Windows Explorer to the folder in which the shortcut for Exchange Management Shell is located. Right-click this shortcut, and then select Properties. In the Properties dialog box, you'll then be able to use the options on the Layout tab to customize the shell.

From the State value you can determine the online status of a monitored resource that is used for transport, connections, or communications. State values you might see include:

- **Online** All the components of the monitored resource are online.
- **Partially Online** Some of the components of the monitored resource are not online.
- **Offline** All the components of the monitored resource are offline.
- **Sidelined** The monitored resource is sidelined and may not be in a fully online state.
- **Functional** The monitored resource is functional but might not be in a fully online state.
- **NotApplicable** An online or offline status is not applicable to this monitored resource.
- **Unavailable** The monitored resource is unavailable.

From the alert value, you can determine the general health status of a monitored resource. Alert values you might see include:

- **Healthy** All the components of the monitored resource are healthy.
- **Degraded** Some of the components of the monitored resource are not healthy.
- **Disabled** The components of the monitored resource have been disabled.

- **Unhealthy** All the components of the monitored resource are not healthy.
- **Sidelined** The monitored resource is sidelined and might not be in a fully healthy state.
- **Repairing** The monitored resource is functional but is recovering from a degraded or unhealthy state.
- **Unavailable** The monitored resource is unavailable.
- **Uninitialized** The monitored resource hasn't been initialized.

If a health set has a status other than healthy or online, you can take a closer look at it by using the `-HealthSet` parameter. List the properties of the health set as shown in this example:

```
Get-ServerHealth -Identity MailServer42 -HealthSet ECP.Proxy | fl
```

You can get a formatted list of every monitor, target resource, and its related health set by entering the following command:

```
Get-ServerHealth localhost | ft name,targetresource,healthsetname
```

The output lists the name of the monitor, the target resource, and the name of the corresponding health set. You can store the output for later reference by redirecting the output to a file. In the following example, `c:\data` is the name of an existing folder, and `Healthset-Reference.txt` is the name of the file to create:

```
(get-serverhealth localhost|ft name,targetresource,healthsetname) >
c:\data\healthset-reference.txt
```

The output will look similar to the following:

Name	TargetResource	HealthSetName
ActiveSyncV2CTPMonitor	ActiveSync	ActiveSync
ActiveSyncCTPMonitor	ActiveSync	ActiveSync
ActiveSyncV2DeepTestMonitor	ActiveSync	ActiveSync.Protocol
ActiveSyncDeepTestMonitor	ActiveSync	ActiveSync.Protocol

Tracking User and Workload Throttling

Whenever you are trying to diagnose and resolve problems with Exchange 2016, you need to keep in mind how user and workload throttling may be affecting performance. All users with mailboxes on servers running Exchange 2016 are subject to user throttling policy.

The default user throttling policy is named the Global Throttling Policy. As the name implies, this policy has global scope and applies throughout the organization. User throttling policies also can have organization and regular scope. If you want to configure user throttling, you should create policies with these scopes rather than modify the Global Throttling Policy.

You can list currently defined user throttling policies by entering `Get-ThrottlingPolicy` at the shell prompt. To create and manage user throttling policies, you can use `New-`

ThrottlingPolicy, Set-ThrottlingPolicy, and Remove-ThrottlingPolicy. You can view throttling policies assigned to users by using Get-ThrottlingPolicyAssociation, and assign user throttling policies to users by using Set-ThrottlingPolicyAssociation.

In addition to user throttling, Exchange Server manages workloads for protocols, features, and services using workload throttling policy. Workloads are automatically throttled to prevent overutilization of system resources and to try to ensure managed resources maintain a healthy state.

Each defined workload has an associated policy and classification. Workload policies are used to enable and configure workloads. Workload classifications set the default priority of the workload. Classifications that can be assigned to workloads include:

- [Urgent](#)
- [Customer Expectation](#)
- [Internal Maintenance](#)
- [Discretionary](#)

You can view the current workload policies and their associated workload classifications by entering Get-WorkloadPolicy at the Shell prompt. To create and manage workload policies, you can use New-WorkloadPolicy, Set-WorkloadPolicy, and Remove-WorkloadPolicy.

Managed resources have health indicators and resource thresholds. Health indicators are used to measure the relative health of the workload in terms of the resources used. Health indicators tracked include:

- [Percent CPU utilization](#)
- [Mailbox database RPC latency](#)
- [Mailbox database replication health](#)
- [Content indexing age of last notification](#)
- [Content indexing retry queue size](#)

Resource thresholds are used to configure usage limits for a system resource. Within each workload classification, one of three thresholds can be assigned: underloaded, overloaded, or critical. As an example:

- [Discretionary workloads are considered underloaded at 70 percent utilization, overloaded at 80 percent utilization, and critical at 100 percent utilization.](#)
- [Internal Maintenance workloads are considered underloaded at 75 percent utilization, overloaded at 85 percent utilization, and critical at 100 percent utilization.](#)
- [Customer Expectation workloads are considered underloaded at 80 percent utilization, overloaded at 90 percent utilization, and critical at 100 percent utilization.](#)

You can view the current resource threshold settings for each workload classification by entering the following command:

```
Get-ResourcePolicy | fl
```

To create and manage resource policies, you can use New-ResourcePolicy, Set-

ResourcePolicy, and Remove-ResourcePolicy. Once you've defined custom workload and resource policies, you can create a policy object based on a particular policy by using New-WorkloadManagementPolicy. You then assign the workload management policy to a server using Set-ExchangeServer with the –WorkloadManagementPolicy and –Server parameters.

Tracking Configuration Changes

As part of your standard operating procedures, you should track changes in the configuration of your Exchange servers. Exchange Management Shell provides the following cmdlets for obtaining detailed information on the current configuration of your Exchange servers:

- **Get-ClientAccessService** Displays configuration details for Client Access services on Mailbox servers.
- **Get-ExchangeServer** Displays the general configuration details for Exchange servers.
- **Get-MailboxServer** Displays configuration details for servers with the Mailbox services.
- **Get-OrganizationConfig** Displays summary information about your Exchange organization.
- **Get-TransportService** Displays configuration details for transport services on servers with the Mailbox or Edge Transport server role.

To get related details for a specific server, you pass the Get-TransportService cmdlet the identity of the server you want to work with, as shown in the following example:

```
Get-TransportService mailserver23 | fl
```

To get related details for all servers, omit the –Identity parameter, as shown in the following example:

```
Get-TransportService | fl
```

When you finalize the configuration of your Exchange servers, you should use these cmdlets to store the configuration details for each server role. To store the configuration details in a file, redirect the output to a file, as shown in the following example:

```
Get-TransportService mailserver23 | fl >  
c:\SavedConfigs\transport2016-0603.txt
```

If you then store the revised configuration, any time you make significant changes you can use this information during troubleshooting to help resolve problems that might be related to configuration changes. To compare two configuration files, you can use the file compare command, fc, at an elevated, administrator command prompt. When you use the following syntax with the fc command, the output is the difference between two files:

```
fc FilePath1 FilePath2
```

where *FilePath1* is the full file path to the first file and *FilePath2* is the full file path to the second file. Here is an example:

```
fc c:\SavedConfigs\transport2016-0603.txt c:\SavedConfigs\  
transport2016-0603.txt
```

Because the files contain configuration details for specific dates, the changes shown in the output represent the configuration changes that you've made to the server.

Testing Service Health, Mail Flow, Replication and More

As part of troubleshooting, you'll often want to determine the status of required services, which can be done using Test-ServiceHealth. The basic syntax is:

```
Test-ServiceHealth [-Server ServerName]
```

Where *ServerName* is the name of the server to test. If you omit a server name, the local server is tested. As shown in the following sample output, Test-ServiceHealth shows you which required services are running and which aren't:

```
Role                : Mailbox Server Role  
RequiredServicesRunning : True  
ServicesRunning      : {IISAdmin, MExchangeADTopology,  
MExchangeDelivery, MExchangeIS, MExchangeMailboxAssistants,  
MExchangeRepl, MExchangeRPC, MExchangeServiceHost,  
MExchangeSubmission, MExchangeThrottling, MExchangeTransportLogSearch,  
W3Svc, WinRM}  
ServicesNotRunning   : {}
```

```
Role                : Client Access Server Role  
RequiredServicesRunning : True  
ServicesRunning      : {IISAdmin, MExchangeADTopology,  
MExchangeMailboxReplication, MExchangeRPC, MExchangeServiceHost, W3Svc,  
WinRM}  
ServicesNotRunning   : {}
```

```
Role                : Unified Messaging Server Role  
RequiredServicesRunning : True  
ServicesRunning      : {IISAdmin, MExchangeADTopology,  
MExchangeServiceHost, MExchangeUM, W3Svc, WinRM}  
ServicesNotRunning   : {}
```

```
Role                : Hub Transport Server Role  
RequiredServicesRunning : True  
ServicesRunning      : {IISAdmin, MExchangeADTopology,  
MExchangeEdgeSync, MExchangeServiceHost, MExchangeTransport,  
MExchangeTransportLogSearch, W3Svc, WinRM}  
ServicesNotRunning   : {}
```

Although Exchange 2016 no longer has separate roles for Mailbox servers, Test-

ServiceHealth continues to list separately the related required services and their status. As part of troubleshooting, you'll often need to test mail flow and replication. If you suspect a problem with mailflow, you can quickly send a test message by using Test-Mailflow. This cmdlet verifies whether mail can be successfully sent from and delivered to the system mailbox as well as whether email is sent between Mailbox servers within a defined latency threshold.

To test mail flow from one mailbox server to another or from one mailbox server to a target mailbox database, you can use the following syntax:

```
Test-MailFlow -Identity OriginatingMailServer [-TargetMailboxServer  
DestinationMailServer | -TargetDatabase DestinationDatabase]
```

In the following example, a test message is sent from MailboxServer34 to MailboxServer26:

```
Test-MailFlow -Identity MailboxServer34 -TargetMailboxServer  
MailboxServer26
```

As shown in this sample, the output of the command tells you whether the message was sent and received successfully:

```
TestMailflowResult : Success  
MessageLatencyTime : 00:00:04.0077377  
IsRemoteTest       : False  
Identity           :  
IsValid            : True  
ObjectState        : New
```

If you suspect a problem with replication, you can quickly determine the status of replication components by using Test-ReplicationHealth. This cmdlet checks the status of all aspects of replication, replay, and availability on a Mailbox server in a Database Availability group. Use Test-ReplicationHealth to help you monitor the status of continuous replication, availability of Active Manager, and the general status of availability components.

The basic syntax is:

```
Test-MailFlow [-Identity MailboxServerId]
```

Such as:

```
Test-MailFlow MailServer19
```

As shown in this sample, the output of the command tells you the status of each replication component on the Mailbox server:

Server	Check	Result	Error
MAILSERVER19	ReplayService	Passed	
MAILSERVER19	ActiveManager	Passed	
MAILSERVER19	TasksRpcListener	Passed	

```
MAILSERVER19 DatabaseRedundancy *FAILED* Failures:...
MAILSERVER19 DatabaseAvailability *FAILED* Failures:...
```

If errors are found, you'll want to get more details by formatting the output in a list, such as:

```
Test-MailFlow MailServer19 | fl server, check*, result, error
```

The error details should help you identify the problem. In this example, the Mailbox database doesn't have enough copies to be fully redundant:

```
Server      : MAILSERVER19
Check       : DatabaseRedundancy
CheckDescription : Verifies that databases have sufficient redundancy. If this check fails, it
means that some databases are at risk of losing data.
Result      : *FAILED*
Error       : Failures:
There were database redundancy check failures for database 'Engineering
Mailbox Database' that may be lowering its redundancy and
putting the database at risk of data loss. Redundancy Count: 1. Expected Redundancy Count: 2.
```

In this example, the Engineering Mailbox Database does not have enough copies for full redundancy. This could be because an administrator forgot to make a passive copy of the database or because a Mailbox server hosting a copy of the database is offline or otherwise unavailable.

Other useful cmdlets for checking the Exchange organization include:

- **Test-ActiveSyncConnectivity** Performs a full synchronization against a specified mailbox to test the configuration of Exchange ActiveSync.
- **Test-ArchiveConnectivity** Verifies archive functionality for a mailbox user.
- **Test-AssistantHealth** Verifies that the Exchange Mailbox Assistant service is running as expected.
- **Test-CalendarConnectivity** Verifies that calendar sharing as part of Outlook Web App is working properly.
- **Test-EcpConnectivity** Verifies that the Exchange Admin Center is running as expected.
- **Test-EdgeSynchronization** Verifies that the subscribed Edge Transport servers have a current and accurate synchronization status.
- **Test-ExchangeSearch** Verifies that Exchange Search is currently enabled and is indexing new email messages in a timely manner.
- **Test-FederationTrust** Verifies that the federation trust is properly configured and functioning as expected.
- **Test-FederationTrustCertificate** Verifies the status of certificates used for federation on all Mailbox servers.
- **Test-ImapConnectivity** Verifies that the IMAP4 service is running as expected.
- **Test-IPAllowListProvider** Verifies the configuration for a specific IP allow list provider.

- **Test-IPBlockListProvider** Verifies the configuration for a specific IP block list provider.
- **Test-IRMConfiguration** Verifies Information Rights Management (IRM) configuration and functionality.
- **Test-MapiConnectivity** Verifies server functionality by logging on to the mailbox that you specify.
- **Test-MRSHealth** Verifies the health of the Microsoft Exchange Mailbox Replication Service.
- **Test-OAuthConnectivity** Verifies that OAuth authentication is working properly.
- **Test-OutlookConnectivity** Verifies end-to-end Microsoft Outlook client connectivity and also tests for Outlook Anywhere (RPC/HTTP) and TCP-based connections.
- **Test-OutlookWebServices** Verifies the Autodiscover service settings for Outlook.
- **Test-OwaConnectivity** Verifies that Outlook Web App is running as expected.
- **Test-PopConnectivity** Verifies that the POP3 service is running as expected.
- **Test-PowerShellConnectivity** Verifies whether Windows PowerShell remoting on the target Mailbox server is functioning correctly.
- **Test-SenderId** Verifies whether a specified IP address is the legitimate sending address for a specified SMTP address.
- **Test-SmtpConnectivity** Verifies SMTP connectivity for a specified server.
- **Test-UMConnectivity** Verifies the operation of a computer that has the Unified Messaging installed.
- **Test-WebServicesConnectivity** Verifies the functionality of Exchange Web Services.

Diagnosing and Resolving Problems

As discussed previously in this chapter in the "Troubleshooting Essentials" section, you can use Get-ServerHealth to list monitors, target resources, and corresponding health sets. Knowing which monitor, target resource, and health set you want to work with is important for troubleshooting. To diagnose and resolve problems, you often need to work backward from the reported problem to the source of the problem, as shown here:

1. Find recovery actions.
2. Trace recovery actions to their responder.
3. Use the responses logged by a responder to find the related monitor.
4. Find the probes for a monitor.
5. Locate the error messages being logged by probes.
6. Verify probe errors still exist.

The sections that follow examine the related procedures.

Identifying Recovery Actions

During recovery, the responder engine uses responders to take appropriate recovery actions, based on the type of alert and the affected target resource. Whenever a responder takes a recovery action, it logs related events in the Microsoft.Exchange.ManagedAvailability/RecoveryActionResults event log. An entry with an event ID of 500 indicates that a recovery action has started. An entry with an event ID of 501 indicates that the recovery action was completed.

Although you can view the events in Event Viewer, you can also view them at the Shell prompt. To collect the events in the RecoveryActionResults event log so you can process them, enter the following commands:

```
$Results = Get-WinEvent -ComputerName ServerName
-LogName Microsoft-Exchange-ManagedAvailability/RecoveryActionResults
```

```
$ResultsXML = ($Results | ForEach-object
-Process {[xml]$_}.event.userData.eventXml
```

Where ServerName is the name of the Mailbox server that you want to work with. The first command collects the events. The second command formats the event entries so that they are easier to work with. These commands can be combined and shortened to:

```
$ResultsXML = (Get-WinEvent -ComputerName ServerName -LogName
Microsoft-Exchange-ManagedAvailability/RecoveryActionResults |
% {[xml]$_}.event.userData.eventXml
```

Next, you need to identify a response that you want to look at more closely. If you want to review corrective actions taken by Managed Availability, you'd look for events that occurred today and completed successfully. The following example parses the

previously collected event data and looks for events from 2016-08-07 that have a successful result:

```
$ResultsXML | Where-Object {$_.Result -eq "Succeeded" -and $_.EndTime -like "2016-08-07*"} | ft -AutoSize StartTime,RequestorName
```

As shown in this example, you also could look for events that occurred but where the responder failed to correct the issue:

```
$ResultsXML | Where-Object {$_.Result -eq "Failed" -and $_.EndTime -like "2016-08-07*"} | ft -AutoSize StartTime,RequestorName
```

With either approach, you'll then get a list of issues by start time and requestor name, such as:

StartTime	RequestorName
2016-08-07t21:00:10.1008312Z	SearchLocalCopyStatusRestartSearchService
2016-08-07t21:00:06.1162578Z	RWSPProxyTestRecycleAppPool
2016-08-07t21:00:00.4597184Z	ClusterEndpointRestart
2016-08-07t20:59:36.1601996Z	RWSPProxyTestRecycleAppPool
2016-08-07t20:57:17.8657794Z	OutlookSelfTestRestart
2016-08-07t20:58:03.7958299Z	RWSPProxyTestRecycleAppPool
2016-08-07t20:55:24.6591276Z	ServiceHealthActiveManagerRestartService
2016-08-07t20:57:11.2223574Z	ClusterEndpointRestart
2016-08-07t20:55:06.9326525Z	OutlookSelfTestRestart
2016-08-07t20:57:02.6438007Z	RWSPProxyTestRecycleAppPool
2016-08-07t20:54:34.5391633Z	OutlookMailboxDeepTestRestart
2016-08-07t20:56:32.4360908Z	RWSPProxyTestRecycleAppPool
2016-08-07t20:54:41.4926429Z	ClusterEndpointRestart
2016-08-07t20:53:34.1596832Z	ActiveDirectoryConnectivityRestart
2016-08-07t20:52:11.0579430Z	ClusterEndpointRestart

In this example, the value in the RequestorName column is the responder that took the action. To examine the properties of a recovery action, run a query for a specific responder, such as:

```
$ResultsXML | Where-Object {$_.Result -eq "Failed" -and $_.EndTime -like "2013*" -and $_.RequestorName -eq "OutlookSelfTestRestart"} | fl
```

The output includes the details logged for events in which the recovery action initiated by the OutlookSelfTestRestart responder failed. Each entry will look similar to the following:

```
auto-ns2      : http://schemas.microsoft.com/win/2004/08/events
xmlns         : myNs
Id            : RestartService
InstanceId    : 130629.015717.86577.001
ResourceName  : MExchangeRPC
StartTime     : 2016-08-07T20:57:17.8657794Z
EndTime      : 2016-08-07T20:59:19.4994266Z
```

State : Finished
Result : Failed
RequestorName : OutlookSelfTestRestart
ExceptionName : TimeoutException
ExceptionMessage : System error.
Context : [null]
CustomArg1 : [null]
CustomArg2 : [null]
CustomArg3 : [null]
LamProcessStartTime : 8/07/2016 1:12:28 PM

Although the responder name and details will often help you identify the type of problem that occurred, you can keep working toward the exact problem that occurred by finding the monitor that triggered the responder.

Identifying Responders

Whenever the Health Manager service starts, it logs related events in the Microsoft.Exchange.ActiveMonitoring/ResponderDefinition event log that you can use to get properties of responders. To collect the events in the ResponderDefinition event log so that you can process them, enter the following command:

```
$Responders = (Get-WinEvent -ComputerName ServerName -LogName  
Microsoft-Exchange-ActiveMonitoring/ResponderDefinition | %  
{[xml]$_}.toXml()).event.userData.eventXml
```

Where **ServerName** is the name of the Mailbox server with which you want to work. If you examine the definition of a responder, the AlertMask property will identify the monitor associated with the responder. Thus, one way to display the required information is to look for the responder and list the responder name and the associated alert mask in the output as shown in this example:

```
$Responders | ? {$_ .Name -eq "OutlookSelfTestRestart"} |  
ft name, alertmask
```

As the output will then be similar to the following:

Name	AlertMask
-----	-----
OutlookSelfTestRestart	OutlookSelfTestMonitor
OutlookSelfTestRestart	OutlookSelfTestMonitor

You'll know the related monitor is named OutlookSelfTestMonitor. Before examining the related monitor, you might want to display the full details for the responder to help you understand exactly how the responder works. To display the full details for a responder, simply list its properties in a formatted list as shown in this example:

```
$Responders | ? {$_ .Name -eq "OutlookSelfTestRestart"} | fl
```

During recovery, the responder engine uses responders to take appropriate recovery actions based on the alert type and the affected target resource. The wait interval

specifies the minimum amount of time a responder must wait before running again. As shown in this partial output, the definition details can help you learn more about the responder:

```
Id : 452
AssemblyPath : C:\Program Files\Microsoft\Exchange
Server\V15\Bin\Microsoft.Exchange.Monitoring.ActiveMonitoring
.Local.Components.dll
TypeName : Microsoft.Exchange.Monitoring
.ActiveMonitoring.Responders.ResetIISAppPoolResponder
Name : OutlookSelfTestRestart
WorkItemVersion : [null]
ServiceName : Outlook.Protocol
DeploymentId : 0
ExecutionLocation : [null]
CreatedTime : 2016-08-07T20:02:32.2527661Z
Enabled : 1
TargetResource : MSExchangeRpcProxyAppPool
RecurrenceIntervalSeconds : 0
TimeoutSeconds : 300
StartTime : 2016-08-07T20:02:32.2527661Z
UpdateTime : 2016-08-07T17:55:07.9754209Z
MaxRetryAttempts : 3
ExtensionAttributes : <ExtensionAttributes
AppPoolName="MSExchangeRpcProxyAppPool" MinimumSecondsBetweenRestarts="300"
MaximumAllowedRestartsInAnHour="3" MaximumAllowedRestartsInADay="-1"
DumpOnRestart="FullDump" DumpPath="C:\Program Files\Microsoft\Exchange
Server\V15\Dumps" MinimumFreeDiskPercent="15" MaximumDumpsPerDay="9"
MaximumDumpDurationInSeconds="180" />
AlertMask : OutlookSelfTestMonitor
WaitIntervalSeconds : 30
MinimumSecondsBetweenEscalates : 0
NotificationServiceClass : 0
AlwaysEscalateOnMonitorChanges : 0
```

Identifying Monitors

Monitor definitions are written in the Microsoft.Exchange.ActiveMonitoring/MonitorDefinition event log. If you examine the properties of events, you can learn more about monitors and learn their related probes. To collect the events in the MonitorDefinition event log so that you can process them, enter the following command:

```
$Monitors = (Get-WinEvent -ComputerName ServerName -LogName
Microsoft-Exchange-ActiveMonitoring/MonitorDefinition | %
{[xml]$_}.event.userData.eventXml
```

Where **ServerName** is the name of the Mailbox server with which you want to work. If you examine the definition of a monitor, the **SampleMask** property will identify the

probes associated with the monitor. List the monitor name and the associated sample mask in the output as shown in this example:

```
$Monitors | ? {$_ .Name -eq "OutlookSelfTestMonitor"} |  
ft name, samplemask
```

The output will then be similar to the following:

```
Name                AlertMask  
----                -  
OutlookSelfTestMonitor      OutlookSelfTestProbe
```

As shown in the output, probes related to this monitor have the top-level identifier: OutlookSelfTestProbe. To display the full details for a monitor, simply list its properties in a formatted list as shown in this example:

```
$Monitors | ? {$_ .Name -eq "OutlookSelfTestMonitor"} | fl
```

During detection, the monitor engine uses monitors to analyze the sampled data. Whether a monitor issues an alert depends on the state of the target resource. As shown in this partial output, the monitor details provide a lot of information, including the exact definition of each transition state for the monitor:

```
Id                : 339  
AssemblyPath      : C:\Program Files\Microsoft\Exchange  
Server\V15\Bin\Microsoft.Exchange.Monitoring.ActiveMonitoring.Local.  
Components.dll  
TypeName          : Microsoft.Exchange.Monitoring.  
ActiveMonitoring .ActiveMonitoring.Monitors  
.OverallConsecutiveProbeFailuresMonitor  
Name              : OutlookSelfTestMonitor  
WorkItemVersion   : [null]  
ServiceName       : Outlook.Protocol  
DeploymentId      : 0  
ExecutionLocation : [null]  
CreatedTime       : 2016-08-07T20:02:32.2215111Z  
Enabled           : 1  
RecurrenceIntervalSeconds : 0  
TimeoutSeconds    : 30  
StartTime         : 2016-08-07T20:02:32.2215111Z  
UpdateTime        : 2016-08-07T19:59:57.2971492Z  
MaxRetryAttempts  : 0  
ExtensionAttributes : [null]  
SampleMask        : OutlookSelfTestProbe  
MonitoringIntervalSeconds : 300  
MinimumErrorCount : 0  
MonitoringThreshold : 2  
SecondaryMonitoringThreshold : 0  
ServicePriority    : 0  
ServiceSeverity    : 0  
IsHaImpacting     : 1
```

```

CreatedById                : 0
InsufficientSamplesIntervalSeconds : 28800
StateAttribute1Mask        : [null]
FailureCategoryMask        : 0
ComponentName              : ServiceComponents/Outlook.Protocol/Critical
StateTransitionsXml        : <StateTransitions>
<Transition ToState="Degraded" TimeoutInSeconds="0" />
<Transition ToState="Degraded1" TimeoutInSeconds="10" />
<Transition ToState="Degraded2" TimeoutInSeconds="240" />
<Transition ToState="Unhealthy" TimeoutInSeconds="300" />
<Transition ToState="Unhealthy1" TimeoutInSeconds="600" />
<Transition ToState="Unrecoverable" TimeoutInSeconds="1200" />
</StateTransitions>
Version                    : 65536

```

Identifying Probes

To identify the probes associated with the OutlookSelfTestProbe identifier, you need to examine the probe definitions. Probe definitions are written in the Microsoft.Exchange.ActiveMonitoring/ProbeDefinition event log. If you examine the properties of events, you can learn more about each probe. To collect the events in the ProbeDefinition event log so that you can process them, enter the following command:

```

$Probes = (Get-WinEvent -ComputerName ServerName -LogName
Microsoft-Exchange-ActiveMonitoring/ProbeDefinition | %
{[xml]$_}.toXml()).event.userData.eventXml

```

Where **ServerName** is the name of the Mailbox server with which you want to work. Next, examine the associated probes to learn more about them as shown in this example:

```

$Probes | ? {$_.Name -eq "OutlookSelfTestProbe"} | fl

```

The output will then list the definition of each associated probe. Although many monitors have many associated probes, the OutlookSelfTestMonitor has only one associated probe. In this partial sample of the output, note the recurrence interval, timeout, and max retry values for this probe:

```

Id                : 106
AssemblyPath      : C:\Program Files\Microsoft\Exchange
Server\V15\Bin\Microsoft.Exchange.Monitoring.ActiveMonitoring
.Local.Components.dll
TypeName          : Microsoft.Exchange.Monitoring.ActiveMonitoring
.RpcClientAccess.LocalRpcProbe+SelfTest
Name              : OutlookSelfTestProbe
WorkItemVersion   : [null]
ServiceName       : Outlook.Protocol
DeploymentId       : 0
ExecutionLocation : [null]
CreatedTime       : 2016-08-07T20:02:32.2058880Z
Enabled           : 1

```

```

RecurrenceIntervalSeconds : 10
TimeoutSeconds           : 8
StartTime                : 2016-08-07T20:02:41.2215111Z
UpdateTime               : 2016-08-07T19:59:57.2190196Z
MaxRetryAttempts        : 0
ExtensionAttributes     : <ExtensionAttributes AccountLegacyDN="
/o=First Organization/ou=Monitoring Mailboxes/cn=Recipients
/cn=HealthMailbox3d899a319e1e4c019f5362ead47f0185"
PersonalizedServerName="278c17fc-8adc-49d7-affa-90f0ea7679b6@
imaginedlands.com" StartupNotificationId="MSExchangeRPC"
StartupNotificationMaxStartWaitInSeconds="12
/>
CreatedById             : 0
Account                 : <r at="Kerberos" ln="IMAGINEDLANDS\SM_fef8fb0aaba040c19">
<s>S-1-5-21-1487214957-3235876329-
1606252878-1151</s><s a="7" t="1">S-1-5-21-1487214957-3235876329-
1606252878-513</s><s a="7" t="1">S-1-1-0</s><s a="7" t="1">S-1-5-2</s>
<s a="7" t="1">S-1-5-11</s><s a="7" t="1">S-1-5-15</s>
<s a="3221225479" t="1">S-1-5-5-0-8194354</s><s a="7" t="1">
S-1-18-2</s></r>
AccountDisplayName      : HealthMailbox3d899a319e1e4c019f5362ead47f0185
Endpoint               : MailServer21.imagedlands.com
SecondaryAccount       : [null]
SecondaryAccountDisplayName : [null]
SecondaryEndpoint      : MailServer21.imagedlands.com
ExtensionEndpoints     : [null]
Version                : 65536
ExecutionType          : 0

```

During sampling, the probe engine runs probes against target resources. How often a probe runs depends on its recurrence interval. How long a probe waits before reporting failure depends on its timeout value. Also listed in the output is the system account under which the probe runs and the authentication method used for that account.

Viewing Error Messages for Probes

Once you know which probes are associated with the issue you are tracking, you can get the error messages for the probes. Probe results are written in the Microsoft.Exchange.ActiveMonitoring/ProbeResult event log. As this log is quite extensive, you want to filter the logs for the exact information you are seeking.

Properties for related events include:

- **ServiceName** Identifies the related health set.
- **ResultName** Identifies the name of the probe. When there are multiple probes for a monitor the name includes the monitor's sample mask and the resource it verifies.
- **Error** Lists the error returned by this probe, if it failed.
- **Exception** Lists the call stack of the error, if it failed.
- **ResultType** Lists an integer value that indicates the result type: 1 for timeout, 2 for

- poisoned, 3 for succeeded, 4 for failed, 5 for quarantined, and 6 for rejected.
- **ExecutionStartTime** Lists when the probe started.
- **ExecutionEndTime** Lists when the probe completed.
- **ExecutionContext** Provides additional information about the probe's execution context.
- **FailureContext** Provides additional information about the probe's failure.

Knowing this, you can collect the events in the ProbeResult event log and filter them. In this example, you look for failure results related to OutlookSelfTestProbe:

```
$Errors = (Get-WinEvent -ComputerName ServerName -LogName
Microsoft-Exchange-ActiveMonitoring/ProbeResult -FilterXPath
"*[UserData[EventXML[ResultName='OutlookSelfTestProbe']][ResultType='4']]")
|% {[XML]$_}.event.userData.eventXml
```

Where ServerName is the name of the Mailbox server with which you want to work. After you filter the log, you can display the results you want to see, such as:

```
$Errors | select -Property *Time,Result*,Error*,*Context
```

In this example, the output lists the time-, result-, error-, and context-related properties, which will help you identify the exact problem that occurred. Consider the following example:

```
ExecutionStartTime : 2016-08-07T21:24:26.9816420Z
ExecutionEndTime  : 2016-08-07T21:24:27.7508864Z
ResultId          : 644887342
ResultName        : OutlookSelfTestProbe
ResultType        : 4
Error              : The request was aborted: Could not create SSL/TLS
secure channel.
ExecutionContext  : RpcProxy connectivity verification
Task produced output:
- TaskStarted = 8/07/2016 2:24:26 PM
- TaskFinished = 8/07/2016 2:24:27 PM
- Exception = System.Net.WebException: The request
was aborted: Could not create SSL/TLS secure channel.
- ErrorDetails = Status: SecureChannelFailure
    HttpStatusCode:
    HttpStatusCodeDescription:
    ProcessedBody:
        - Latency = 00:00:00.5617493
- RpcProxyUrl = https://mailserver21.
imaginedlands.com:444/rpc/rpcproxy.dll?MailServer21.
imaginedlands.com:6001
    - ResponseStatusCode = <null>
    RpcProxy connectivity verification failed.
FailureContext     : Status: SecureChannelFailure
    HttpStatusCode:
    HttpStatusCodeDescription:
```

ProcessedBody:

As you can see from the output, the probe error details provide a lot of information regarding the exact problem that occurred. In this example, an RPC Proxy error occurred that prevented creation of a secure SSL/TLS channel. If this was a problem preventing access to the server or causing other issues, you would then know that you need to look at related components to continue your troubleshooting. You would look at the RPC, RPC Proxy, SSL and TLS configuration in Internet Information Services (IIS) as well as the related settings in Exchange.

Tracing Probe Errors

Now that you know how to trace a reported problem to its source, let's take a look at additional ways in which you can put this knowledge to use. You view the overall health of a server by using `Get-ServerHealth`. As discussed earlier in this chapter, if a health set has a status other than healthy or online, you can take a closer look at it by using the `-HealthSet` parameter. List the properties of the health set as shown in this example:

```
Get-ServerHealth -Identity MailServer42 -HealthSet FrontEndTransport | fl
```

The `Name` property in the output of `Get-ServerHealth` lists the name of the monitor reporting the health status. Table 30-1 lists the health sets associated with key Exchange features and components.

TABLE 30-1 Health Sets Associated with Key Exchange Features and Components

FEATURE/COMPONENT	RELATED HEALTH SETS
ActiveSync	ActiveSync, ActiveSync.Protocol, ActiveSync.Proxy
Active Directory	AD
Anti-virus	Antimalware, AntiSpam
Autodiscover	Autodiscover, Autodiscover.Protocol, Autodiscover.Proxy
Mailbox databases	Clustering, Database, DataProtection, MailboxMigration, MailboxSpace, MRS, Store
Exchange Admin Center	ECP.Proxy

Exchange Web Services	EWS, EWS.Protocol, EWS.Proxy
Front-End Transport Service	FrontendTransport
Transport Service	HubTransport, MailboxTransport, Transport, TransportSync
Offline Address Book	OAB, OAB.Proxy
Outlook, Outlook Web App	Outlook, Outlook.Proxy, OWA.Protocol, OWA.Protocol.Dep, OWA.Proxy
Unified Messaging	UM.Callrouter, UM.Protocol
User Throttling	UserThrottling

You can quickly identify all the related probes, monitors, and responders for a health set by using `Get-MonitoringItemIdentity`. The basic syntax is:

```
Get-MonitoringItemIdentity -Identity HealthSetName -Server ServerName
```

Where `HealthSetName` identifies the health set to examine and `ServerName` is the name of an Exchange server. In the following example, you list items by type, item name, and target resource:

```
Get-MonitoringItemIdentity -Identity FrontEndTransport -Server mailserv38
| ft itemtype, name, targetresource
```

As shown in the following partial output, each associated probe, monitor, and responder is listed by name:

```
ItemType Name                               TargetResource
-----
Probe FrontendTransportServiceRunning         msexchangefrontendtransport
Probe FrontendTransportRepeatedlyCrashing msexchangefrontendtransport
Monitor FrontendTransportServiceRunningMonitor
Monitor FrontendTransportRepeatedlyCrashingMonitor
Responder FrontendTransportServiceRunningEscalateResponder Transport
Responder FrontendTransportRepeatedlyCrashingResponder Transport
```

If the name of the monitor reporting a status other than online or healthy is `FrontendTransportRepeatedlyCrashingMonitor`, you can analyze the problem by looking at errors for the `FrontendTransportRepeatedlyCrashing` probe. Collect events for this probe from the `ProbeResult` event log and filter them as discussed earlier in “Viewing error messages for probes.” Here is an example:

```
$Errors = (Get-WinEvent -ComputerName ServerName -LogName
Microsoft-Exchange-ActiveMonitoring/ProbeResult -FilterXPath
"*[UserData[EventXML[ResultName='FrontendTransportRepeatedlyCrashing']
[ResultType='4']]") | % {[XML]$_}.event.userData.eventXml
```

Where `ServerName` is the name of the Mailbox server with which you want to work. Remember, the result type can be 1 for timeout, 2 for poisoned, 3 for succeeded, 4 for failed, 5 for quarantined, or 6 for rejected.

After you filter the log, you can display the results you want to see, such as:

```
$Errors | select -Property *Time,Result*,Error*,*Context
```

Before you begin deeper troubleshooting, you might want to rerun the associated probe for the monitor to ensure its still not in a healthy or online state. You can rerun probes by using `Invoke-MonitoringProbe`. The basic syntax is:

```
Invoke-MonitoringProbe HealthSetName\ProbeName -Server ServerName | fl
```

Where `HealthSetName` is the name of the health set with which to work, `ProbeName` is the name of the probe within the specified health set, and `ServerName` is the name of the Exchange server to check, such as:

```
Invoke-MonitoringProbe FrontEndTransport\
FrontendTransportRepeatedlyCrashing -Server MailServer38 | fl
```

As shown in this partial sample of the output, the command returns a lot of information about the test:

```
Server          : MailServer38
MonitorIdentity : FrontEndTransport\FrontendTransportRepeatedlyCrashing
RequestId       : 84dc68cd-c2f8-487f-a5e2-20b43f6f9207
ExecutionStartTime : 6/5/2016 8:20:42 AM
ExecutionEndTime  : 6/5/2016 8:20:42 AM
Error           :
Exception       :
PoisonedCount    : 0
ExecutionId      : 18902819
SampleValue      : 2015
ExecutionContext :
FailureContext   :
ExtensionXml     :
ResultType       : Succeeded
RetryCount       : 0
ResultName       : 84dc68cdc2f8487fa5e220b43f6f9207-
FrontendTransportRepeatedlyCrashing
IsNotified       : False
ResultId         : 1289896134
ServiceName      : InvokeNow
StateAttribute1  : No relevant crash events found for service
```

The `ResultType` value in the output will tell you whether the probe succeeded or failed.

If the probe succeeded, the problem no longer exists. If the probe fails, the problem still exists and you'll need to continue trying to diagnose and resolve it. Step by step procedures for troubleshooting issues with Exchange services are provided in the "Troubleshooting Outlook Web App" section of this chapter and "Maintaining Virtual Directories and Web Applications" section of Chapter 25, "Optimizing Web and Mobile Access."

@techjob

Troubleshooting Outlook Web App

As discussed in “Troubleshooting OWA, ECP, Powershell, and More” in Chapter 16, “Exchange 2016 Administration Essentials,” sometimes users and administrators see a blank page or an error when they try to log on to OWA. This problem and other connection issues, such as those related to ECP, OAB, Autodiscover, and Windows PowerShell, can occur because of a wide variety of configuration issues, including:

- Invalid or missing TCP/IP settings.
- Corrupted or improperly configured virtual directories.
- Missing, expired, invalid, or improperly configured SSL certificates.

You resolve these issues by correcting the configuration problem as discussed in that chapter. Beyond configuration issues, Exchange servers can have connectivity, resource, and service issues. You can use Test-OwaConnectivity to test connectivity to Outlook Web App as part of troubleshooting connectivity; however, this cmdlet is deprecated and will be removed in a future release of Exchange Server.

Checking OWA Health

Exchange 2016 uses Active Monitoring to monitor essential services, connectivity, resources and the overall health of the messaging platform. Active Monitoring is performed by the Microsoft Exchange Health Manager service, which must be running on the Exchange server. As discussed in detail in “Managed Availability Components” in Chapter 18 “Implementing Availability Groups,” Active Monitoring is itself part of the Managed Availability feature.

The overall health of Outlook Web App is tracked by the OWA health set. A health set includes a probe that takes measurements on the server and collects data, and a monitor that uses the collected data to determine whether a resource is healthy. OWA relies on the OwaCtpProbe to measure the health of Outlook Web App and the OwaCtpMonitor to determine the status of Outlook Web App. The OWA health is dependent on Active Directory Domain Services (AD DS) and the Microsoft Exchange Information Store service.

Alerts related to resources are logged in the event logs. You also can manually check the status of resources by using the Get-HealthReport and Get-ServerHealth. Whereas Get-ServerHealth provides the exact state and health of every Exchange resource, monitor, and service, Get-HealthReport returns the state of monitored resources. You can quickly check for unhealthy resources by entering the following command:

```
Get-ServerHealth -Identity ServerID | where ($_.AlertValue -eq 'Unhealthy')
```

Where *ServerID* is the host name or fully-qualified name of the Exchange server to check, such as:

```
Get-ServerHealth -Identity MailServer21.pocketonconsultant.com |
```

where (`$_AlertValue -eq 'Unhealthy'`)

Rather than check all resources and health sets, you can explicitly check the status of the OWA-related health sets by using the following command:

```
Get-ServerHealth ServerID | ?{$_HealthSetName -match "OWA"}
```

Where *ServerID* is the host name or fully-qualified name of the Exchange server to check, such as:

```
Get-ServerHealth MailServer21.pocketonconsultant.com |  
?{$_HealthSetName -match "OWA"}
```

NOTE Here, I've used a filter that looks for values that contain a match for OWA rather than a filter that looks for a value that equals OWA. In this way, you get the status of every OWA related health set rather than just the OWA health set.

Mailbox servers use IIS for front-end services, such as authentication and proxying, as well as back-end processing. You'll find front-end apps for OWA, ECP, PowerShell, OAB, and Autodiscover apps are configured on the Default Web Site. You'll find back-end apps for OWA, ECP, PowerShell, OAB, and Autodiscover are configured on Exchange Back End website.

Understanding Unhealthy Status

If the OWA health set reports an unhealthy status, an issue is present that might prevent users from accessing their mailboxes in Outlook Web App. Such issues include:

- The OWA application pool is not responding on the Mailbox server providing front-end proxy services
- The OWA application pool is not responding on the Mailbox server providing back-end services
- Network issues are preventing the Mailbox server from connecting to other Mailbox servers or a domain controller
- A domain controller or the Microsoft Exchange Information Store service is not responding
- The user's mailbox database is dismounted or otherwise inaccessible
- The credentials for the monitoring account are incorrect

Some of these problems can be resolved automatically by the responder engine, which is another Managed Availability component. When a problem exists with application pools or services on Exchange servers, the responder engine attempts to recover the resource by restarting the application pool or service that is causing the problem. The problem identification and recovery process can take several minutes. If you notice a problem with Outlook Web App that you suspect is related to application pools or services, you can, of course, perform the restart procedures yourself to try to restore access more quickly to Outlook Web App.

OWA.Proxy and OWA.Protocol also are related health sets. OWA.Proxy relies on OwaProxyTestProbe, OWAAnonymousCalendarProblem, and OwaProxyTestMonitor to

track the status of proxy services and calendaring features. OWA.Protocol relies on:

- **OwaSelfTestProbe** and the **OwaSelfTestMonitor**. **OwaSelfTestProbe** performs connectivity tests by sending an HTTP request to *https://localhost:444/owa/exhealth.check* . If the probe gets back a status code of 200 OK, the **MSExchangeOWAAppPool** is responding. This probe doesn't depend on any other Exchange component.
- **OwaDeepTestProbe** and **OwaDeepTestMontor**. **OwaDeepTestProbe** checks each Mailbox database on the server to ensure that mailbox users can log on to the server using Outlook Web App. This probe depends on Active Directory for authentication and the Microsoft Exchange Information Store for mailbox access.

As with the OWA health set, an unhealthy status for OWA.Proxy or OWA.Protocol means an issue exists that might prevent users from accessing their mailboxes in Outlook Web App. The common issues for the OWA.Protocol health set are the same as those for the OWA health set. With OWA.Proxy, common issues may be related to the OWA application pool not responding on the Mailbox server providing front-end proxy services, a domain controller not responding, or the credentials for the monitoring account being incorrect.

Correcting Unhealthy Status

You can diagnose a problem with OWA by using **Get-HealthReport** to check the status of the OWA related health set. If the problem you are experiencing with Outlook Web App isn't a configuration issue, use the following techniques to try to resolve the problem while verifying the issue still exists each time you take a corrective action:

1. Try to isolate the problem to a specific server by running a health check for each server. If OWA.Proxy or OWA.Protocol for a particular server has an Unhealthy status, you've likely isolated the problem and identified the server experiencing the problem and can skip Steps 2 and 3.
2. If you are unable to isolate the problem to a specific server or servers, try to access and log on to Outlook Web App by using the URL for a specific Mailbox server. If this fails, try accessing and logging on to a different Mailbox server to help you verify whether the problem is with a particular server. Remember that the Mailbox server used in the one that contains the mailbox database where the mailbox for the user is stored.
3. Using the Services console, verify that all essential Exchange services are running on the Mailbox servers. If an essential service isn't running, select it, and then click **Start**.
4. Verify network connectivity between the Mailbox servers. One way to do this is to log on to each server and try to ping the other servers. If you correct a connectivity issue, check to see if the OWA issue is resolved.
5. In IIS Manager, connect to the server that's reporting the health issue or otherwise experiencing a problem with OWA. Expand the Sites node and verify that the Default Web Site or Exchange Back End website is running as appropriate. If a

required website isn't running, click Start in the Actions pane to start it. This should resolve the problem.

6. Under Application Pools, verify that the required application pools have been started. If a required application pool hasn't been started, select it, and then click Start in the Actions pane.

The main application pool for OWA is MExchangeOWAAppPool. A single application pool with this name is used for both front-end and back-end services. For calendaring, OWA relies on MExchangeOWACalendarAppPool. A single application pool with this name is used for both front-end and back-end services.

REAL WORLD As your messaging environment grows and usage of Outlook Web App increases, you may find that the basic application pool settings for MExchangeOWAAppPool are insufficient. Specifically, if users are getting an HTTP 503 "Service Unavailable" response when they try to connect to OWA, you may need to increase the queue length so that a greater number of requests can be queued in the application pool. Although slow response times likely can be attributed to connection speed and latency on the network, they might also be because the application pool has to service too many users. If so, you may want to consider configuring the Maximum Worker Processes setting so that multiple worker processes can be used. In both cases, doing so, however, requires that additional system resources (primary memory resources) must be allocated to the application pool.

7. If you suspect an issue with MExchangeOWAAppPool, select MExchangeOWACalendarAppPool, and then click Recycle in the Actions pane to recycle its work processes.
8. If you suspect an issue with MExchangeOWACalendarAppPool, select MExchangeOWACalendarAppPool, and then click Recycle in the Actions pane to recycle its work processes.
9. If the problem isn't resolved yet, restart the website where the problem is occurring or the IIS itself. To restart a website, select the website in IIS Manager and then select Restart in the Actions pane. To restart IIS, select the server node in IIS Manager, and then click Restart in the Actions pane.
10. If the problem still isn't resolved, restart the server. If restarting the server doesn't resolve the problem, you likely have a configuration issue that needs to be resolved.

About the Author



William R. Stanek (<http://www.williamrstanek.com>) has more than 20 years of hands-on experience with advanced programming and development. He is a leading technology expert, an award-winning author, and a pretty-darn-good instructional trainer. Over the years, his practical advice has helped millions of programmers, developers, and network engineers all over the world. In 2013, William celebrated the publication of his 150th book.

William has been involved in the commercial Internet community since 1991. His core business and technology experience comes from more than 11 years of military service. He has substantial experience in developing server technology, encryption, and Internet solutions. He has written many technical white papers and training courses on a wide variety of topics. He frequently serves as a subject matter expert and consultant.

William has an MS with distinction in information systems and a BS in computer science, magna cum laude. He is proud to have served in the Persian Gulf War as a combat crewmember on an electronic warfare aircraft. He flew on numerous combat missions into Iraq and was awarded nine medals for his wartime service, including one of the United States of America's highest-flying honors, the Air Force Distinguished Flying Cross. Currently, he resides in the Pacific Northwest with his wife and children.


William recently rediscovered his love of the great outdoors. When he's not writing, he can be found hiking, biking, backpacking, traveling, or trekking in search of adventure with his family! In his spare time, William writes books for children, including *The Bugville Critters Explore the Solar System* and *The Bugville Critters Go on Vacation*.

Find William on Twitter at <http://www.twitter.com/WilliamStanek> and on Facebook at <http://www.facebook.com/William.Stanek.Author>.

Thanks for purchasing *Exchange Server 2016: The Administrator's Reference*.



Please look for these and other William Stanek books at your favorite bookstores and online.



Thank you for purchasing this book. If you found this book to be useful, helpful or informative, raise your voice and support William's work by sharing about this book online.

Unsure how to share? Here are some tips:

- [Blog about the book](#)
- [Write a review at your favorite online store](#)
- [Post about the book on Facebook or elsewhere](#)
- [Tweet about the book](#)

Stay in touch with William on Facebook and Twitter!

Table of Contents

How to Use This Guide	19
Print Readers	19
Digital Book Readers	19
Support Information	19
Conventions & Features	20
Share & Stay in Touch	20
Chapter 1. Welcome to Exchange 2016	22
Getting Started with Exchange Admin Center	23
Navigating Exchange Admin Center Options	23
Accessing Exchange Admin Center	25
Authenticating and Proxying Connections	27
Getting Started with Exchange Management Shell	28
Running and Using Cmdlets	28
Running and Using Other Commands and Utilities	29
Using Cmdlet Parameters and Errors	29
Using Cmdlet Aliases	30
Working with Exchange Management Shell	31
Starting Exchange Management Shell	31
Using Exchange Cmdlets	32
Working with Object Sets and Redirecting Output	33
Chapter 2. Working with Exchange Online	34
Getting Started with Exchange Online	35
Navigating Exchange Online Services	35
Understanding Office 365 Licensing	39
Using Windows PowerShell with Exchange Online	41
Getting Started with Windows PowerShell	41
Understanding the Default Working Environment	42
Learning About Cmdlets and Functions	44
Connecting to Exchange Online Using PowerShell	45
Exploring How the Shell Uses Remote Sessions	45
Establishing Remote Sessions	46
Using an Interactive Remote Session	46
Creating and Importing a Remote Session	48

Connecting to Windows Azure	50
Cmdlets for Windows Azure Active Directory	51
Working with Exchange Online Cmdlets	54
Cmdlets Specific to Exchange Online	54
Working with Exchange Online Cmdlets	56
Chapter 3. Getting Started with Users and Contacts	59
Working with Users and Contacts	60
How Email Routing Works: The Essentials	62
Managing Recipients: The Fundamentals	64
Finding Existing Mailboxes, Contacts, And Groups	69
Finding Synced, Unlicensed, Inactive, and Blocked Users	73
Chapter 4. Managing Users	74
Creating Mailbox-Enabled and Mail-Enabled User Accounts	75
Working with Logon Names and Passwords	75
Mail-Enabling New User Accounts	75
Mail-Enabling Existing User Accounts	80
Managing Mail-Enabled User Accounts	82
Creating Domain User Accounts with Mailboxes	83
Creating Online User Accounts with Mailboxes	89
Adding Mailboxes to Existing Domain User Accounts	92
Setting or Changing the Common Name and Logon Name for Domain User Accounts	96
Setting or Changing Contact Information for User Accounts	97
Changing Logon ID or Logon Domain for Online Users	99
Changing a User's Exchange Server Alias and Display Name	100
Adding, Changing, and Removing Email and Other Addresses	102
Setting a Default Reply Address for a User Account	104
Changing A User's Web, Wireless Service, And Protocol Options	105
Requiring Domain User Accounts to Change Passwords	107
Deleting Mailboxes from User Accounts	108
Deleting User Accounts and Their Mailboxes	110
Chapter 5. Managing Contacts	113
Creating Mail-Enabled Contacts	114

Setting or Changing a Contact's Name and Alias	117
Setting Additional Directory Information for Contacts	118
Changing Email Addresses Associated with Contacts	120
Disabling Contacts and Removing Exchange Attributes	123
Deleting Contacts	125
Chapter 6. Adding Special-Purpose Mailboxes	126
Using Room and Equipment Mailboxes	127
Adding Room Mailboxes	132
Adding Equipment Mailboxes	135
Adding Linked Mailboxes	138
Working with Archive Mailboxes	141
Adding In-Place Archives	141
Adding Online Archives	144
Managing Archive Settings	145
Adding Arbitration Mailboxes	150
Adding Discovery Mailboxes	151
Adding Shared Mailboxes	153
Adding Public Folder Mailboxes	156
Chapter 7. Managing Mailboxes	160
Managing Mailboxes: The Essentials	161
Viewing Current Mailbox Size, Message Count, and Last Logon	161
Configuring Apps for Mailboxes	163
Hiding Mailboxes from Address Lists	165
Defining Custom Mailbox Attributes for Address Lists	166
Restoring On-Premises Users and Mailboxes	167
Restoring Online Users and Mailboxes	170
Repairing Mailboxes	172
Moving Mailboxes	174
Importing and Exporting Mail Data	174
Performing On-Premises Mailboxes Moves and Migrations	175
Performing On-Premises Mailbox Moves	179
Moving Mailboxes Within a Single Forest	180
Moving Mailboxes Between Forests	184
Managing Delivery Restrictions, Permissions, and Storage Limits	189

Setting Message Size Restrictions for Contacts	189
Setting Message Size Restrictions on Delivery to and from Individual Mailboxes	189
Setting Send and Receive Restrictions for Contacts	191
Setting Message Send and Receive Restrictions on Individual Mailboxes	191
Permitting Others to Access a Mailbox	193
Forwarding Email to a New Address	196
Setting Storage Restrictions on Mailbox and Archives	197
Setting Deleted Item Retention Time on Individual Mailboxes	200
Chapter 8. Managing Groups	202
Using Security and Distribution Groups	203
Group Types, Scope, And Identifiers	203
When to Use Security and Standard Distribution Groups	204
When to Use Dynamic Distribution Groups	205
Working with Security and Standard Distribution Groups	207
Group Naming Policy	207
Understanding Group Naming Policy	207
Defining Group Naming Policy for Your Organization	208
Defining Blocked Words in Group Naming Policy	209
Creating Security and Standard Distribution Groups	210
Creating a New Group	211
Mail-Enabling Universal Security Groups	214
Assigning and Removing Membership for Individual Users, Groups, and Contacts	215
Adding and Removing Managers	217
Configuring Member Restrictions and Moderation	218
Working with Dynamic Distribution Groups	221
Creating Dynamic Distribution Groups	221
Changing Query Filters and Filter Conditions	224
Designating an Expansion Server	225
Modifying Dynamic Distribution Groups Using Cmdlets	226
Previewing Dynamic Distribution Group Membership	228
Other Essential Tasks for Managing Groups	230
Changing a Group's Name Information	230
Changing, Adding, or Deleting a Group's Email Addresses	230
Hiding Groups from Exchange Address Lists	232
Setting Usage Restrictions on Groups	233

Creating Moderated Groups	234
Deleting Groups	236
Chapter 9. Managing Addresses Online and Offline	238
Managing Online Address Lists	239
Using Default Address Lists	239
Using Address Book Policies	240
Creating and Applying New Address Lists	242
Updating Address List Configuration and Membership Throughout the Domain	246
Previewing and Editing Address Lists	246
Configuring Clients to Use Address Lists	248
Renaming and Deleting Address Lists	249
Managing Offline Address Books	250
Creating Offline Address Books	250
Configuring Clients to Use an Offline Address Book	251
Setting the Default Offline Address Book	252
Changing Offline Address Book Properties	252
Designating OAB Generation Servers and Schedules	253
Rebuilding the OAB	255
Deleting Offline Address Books	255
Chapter 10. Configuring Exchange Clients	257
Mastering Outlook Web App essentials	258
Getting started with Outlook Web App	258
Connecting to Mailboxes and Public Folder Data Over the Web	259
Working with Outlook Web App	262
Enabling and Disabling Web Access for Users	265
Configuring Mail Support for Outlook	267
Understanding Address Lists, Offline Address Books, and Autodiscover	267
Configuring Outlook for the First Time	269
First-Time Configuration: Connecting to Exchange Server	271
First-Time Configuration: Connecting to Internet Email Servers	273
Configuring Outlook for Exchange	275
Adding Internet Mail Accounts to Outlook	276
Repairing and Changing Outlook Mail Accounts	276
Leaving Mail on the Server with POP3	280
Checking Private and Public Folders with IMAP4 and UNIX Mail	282

Servers	
Managing the Exchange Configuration in Outlook	284
Managing Delivery and Processing Email Messages	284
Using Server Mailboxes	284
Using Personal Folders	284
Repairing .pst data files	287
Repairing .ost data files	289
Accessing Multiple Exchange Mailboxes	290
Logging on to Exchange as the Mailbox Owner	291
Delegating Mailbox Access	291
Opening Additional Exchange Mailboxes	293
Granting Permission to Access Folders Without Delegating Access	294
Using Mail Profiles to Customize the Mail Environment	297
Creating, Copying, and Removing Mail Profiles	297
Selecting a Specific Profile to use on Startup	298
Chapter 11. Customizing & Troubleshooting the Exchange Shell	301
Running and using the Exchange Management Shell	302
Managing the PowerShell Application	309
Customizing Exchange Management Shell	311
Performing One-to-Many Remote Management	313
Using a Manual Remote Shell to Work with Exchange	316
Preparing to Use the Remote Shell	316
Connecting Manually to Exchange 2016 Servers	318
Connecting Manually to Exchange Online	319
Managing Remote Sessions	320
Troubleshooting Exchange Management Shell	321
Chapter 12. Customizing & Configuring Exchange Security	323
Configuring Standard Exchange Permissions	324
Assigning Permissions: Exchange Server and Online	324
Understanding Exchange Management Groups	325
Assigning Management Permissions	330
Understanding Advanced Exchange Server Permissions	333
Assigning Advanced Exchange Server Permissions	336
Configuring Role-Based Permissions for Exchange	338
Understanding Role-Based Permissions	338
Working with Role Groups	344

Managing Role Group Members	349
Assigning Roles Directly or Via Policy	350
Configuring Account Management Permissions	355
Managing Advanced Permissions	358
Adding Custom Roles	358
Adding Custom Role Scopes	360
Adding Custom Role Entries	362
Working with Shared and Split Permissions	366
Using Shared Permissions	366
Using Split Permissions	367
Chapter 13. Implementing Exchange Services	370
Selecting Hardware for Exchange 2016	373
Navigating Exchange 2016 Editions	376
Using Exchange 2016 with Windows Server	382
Services for Exchange Server	382
Exchange Server Authentication and Security	383
Exchange Server Security Groups	384
Using Exchange 2016 with Active Directory	386
Understanding How Exchange Stores Information	386
Understanding How Exchange Routes Messages	386
Additional Tools and Options	388
Chapter 14. Preparing for Exchange 2016	392
Designing the Exchange Server Organization	394
Planning for High Availability	396
Planning Exchange Databases and Storage	398
Planning for Client Access	401
Planning to Support Transport Services	404
Planning for Unified Messaging	406
Integrating Exchange with Active Directory	407
How Mailbox Servers use Active Directory	407
How Edge Transports use Active Directory	408
Integrating Exchange 2016 Into Existing Organizations	410
Coexistence and Active Directory	410
Configuring Exchange 2016 for Coexistence	413

Setting the Default Offline Address Book	414
Moving to Exchange Server 2016	415
Chapter 15. Deploying Exchange Server 2016	417
Installing New Exchange Servers	417
Installing Exchange Server	419
Verifying and Completing the Installation	427
Uninstalling Exchange 2016	430
Using Cumulative Updates	432
What's in Cumulative Updates?	432
How Are Cumulative Updates Applied?	433
How Do I Track Exchange Version Numbers?	433
Installing Cumulative Updates and Service Packs	435
Preparing to Install a Cumulative Update or Service Pack	435
Installing a Cumulative Update or Service Pack	436
Chapter 16. Exchange 2016 Administration Essentials	439
Working with Exchange Admin Center	440
Accessing Exchange Admin Center	440
Working with Exchange Server Certificates	441
Configuring Exchange Admin Center	444
Bypassing Exchange Admin Center and Troubleshooting	448
Understanding Remote Execution in Exchange Admin Center	448
Bypassing Exchange Admin Center and Exchange Management Shell	449
Troubleshooting OWA, ECP, Powershell, and More	450
Resolving SSL Certificate Issues	452
Resolving OWA, ECP, or Other Virtual Directory Issues	453
Validating Exchange Server Licensing	455
Using and Managing Exchange Services	458
Working with Exchange Services	458
Checking Required Services	459
Maintaining Exchange Services	460
Configuring Service Startup	460
Configuring Service Recovery	461
Customizing Remote Management Services	462
Chapter 17. Managing Exchange Organizations	467
Navigating Exchange 2016 Organizations	468

Organizational Architecture	468
Front End Transport	469
Back End Transport	471
Understanding Exchange Routing	475
Routing Boundaries	475
IP Site Links	476
Cross-Premises Routing	477
Understanding Data Storage in Exchange Server 2016	480
Working with the Active Directory Data Store	480
Using Multimaster Replication	480
Using Global Catalogs	480
Using Dedicated Expansion Servers	481
Navigating the Exchange Information Store	482
Data Storage Components	482
The Managed Store	483
Exchange Server Data Files	484
Data Storage in Exchange Databases	485
Exchange Server Message Queues	486
Chapter 18. Implementing Availability Groups	490
Building Blocks for High Availability	491
The Extensible Storage Engine	491
The High Availability Framework	493
Cluster Components	498
Active Manager Framework	499
Managed Availability Components	500
Creating and Managing Database Availability Groups	503
Preparing for DAGs	503
Creating Database Availability Groups	512
Managing Availability Group Membership	515
Managing Database Availability Group Networks	518
Changing Availability Group Network Settings	522
Configuring Database Availability Group Properties	525
Removing Servers from a Database Availability Group	527
Removing Database Availability Groups	528
Maintaining Database Availability Groups	529
Switching Over Servers and Databases	529
Checking Continuous Replication Status	532

Restoring Operations After a DAG Member Failure	533
Chapter 19. Configuring Exchange Databases	536
Getting Started with Active Mailbox Databases	537
Planning for Mailbox Databases	537
Preparing for Automatic Reseed	538
Creating and Managing Active Databases	541
Creating Mailbox Databases	541
Setting the Default Offline Address Book	545
Setting Mailbox Database Limits and Deletion Retention	546
Recovering Deleted Mailboxes	549
Recovering Deleted Items from Servers	551
Creating and Managing Database Copies	553
Creating Mailbox Database Copies	553
Configuring Database Copies	556
Suspending and Resuming Replication	558
Activating Lagged Database Copies	559
Updating Mailbox Database Copies	561
Monitoring Database Replication Status	565
Removing Database Copies	567
Maintaining Mailbox Databases	569
Checking Database Status	570
Setting the Maintenance Interval	571
Renaming Databases	573
Mounting and Dismounting Databases	573
Configuring Automatic Mounting	575
Moving Databases	575
Deleting Databases	578
Managing Content Indexing	580
Indexing Essentials	580
Maintaining Exchange Store Search	580
Resolving Indexing Issues	582
Chapter 20. Managing SMTP Connectors	583
Send and Receive Connectors: The Essentials	584
Understanding Send and Receive Connectors	584
Routing Messages within Sites	585

Routing Messages Across Site Links	588
Managing Send Connectors	590
Creating Send Connectors	590
Viewing and Managing Send Connectors	600
Configuring Send Connector DNS Lookups	603
Setting Send Connector Limits	604
Managing Receive Connectors	607
Creating Receive Connectors	608
Configuring Receive Connectors	614
Creating Connectors with Exchange Online	618
Chapter 21. Configuring Transport Services	621
Optimizing Transport Limits	623
Setting Organizational Transport Limits	624
Setting Connector Transport Limits	625
Setting Server Transport Limits	626
Setting Exchange Activesync Limits	626
Setting Exchange Web Services Limits	628
Setting Outlook Web App Limits	629
Managing Message Transport	632
Configuring the Postmaster Address and Mailbox	632
Configuring Shadow Redundancy	634
Configuring Safety Net	638
Enabling Anti-Spam Features	639
Subscribing Edge Transport Servers	642
Creating an Edge Subscription	642
Getting Edge Subscription Details	643
Synchronizing Edge Subscriptions	644
Verifying Edge Subscriptions	645
Removing Edge Subscriptions	646
Chapter 22. Maintaining Mail Flow	647
Managing Message Routing and Delivery	648
Understanding Message Pickup and Replay	648
Configuring and Moving the Pickup and Replay Directories	649
Changing the Message Processing Speed	650
Configuring Messaging Limits for the Pickup Directory	651
Configuring Message Throttling	652

Understanding Back Pressure	653
Creating and Managing Accepted Domains	655
Understanding SMTP Domains	655
Viewing Accepted Domains	656
Creating Accepted Domains	657
Changing The Accepted Domain Type and Identifier	659
Removing Accepted Domains	660
Creating and Managing Remote Domains	662
Viewing Remote Domains	662
Creating Remote Domains	662
Configuring Messaging Options for Remote Domains	663
Removing Remote Domains	664
Chapter 23. Implementing Exchange Policies and Rules	666
Creating and Managing Email Address Policies	667
Working with Email Address Policies	667
Creating Email Address Policies	668
Editing and Applying Email Address Policies	672
Removing Email Address Policies	673
Configuring Journal Rules	675
Setting The NDR Journaling Mailbox	675
Creating Journal Rules	676
Managing Journal Rules	677
Configuring Transport Rules	679
Creating Transport Rules	680
Managing Transport Rules	681
Chapter 24. Filtering Spam	683
Filtering Spam by Sender	684
Filtering Spam by Recipient	687
Filtering Connections with IP Block Lists	689
Applying IP Block Lists	689
Configuring Block List Providers	691
Specifying Custom Error Messages	692
Defining Block Lists	693
Using Connection Filter Exceptions	693
Using Global Allowed Lists	693
Using Global Block Lists	694

Preventing Internal Servers from Being Filtered	696
Chapter 25. Optimizing Web and Mobile Access	697
Navigating IIS Essentials for Exchange Server	698
Understanding Mobile Access via IIS	698
Maintaining Virtual Directories and Web Applications	699
Starting, Stopping, and Restarting Websites	703
@techjob	
Configuring Outlook Web App Features	705
Managing Segmentation Features	705
Managing Outlook Web App Policies	707
Managing Bindings, Connections and Authentication	710
Optimizing the Mobile Access Websites	710
Enabling SSL on Websites	712
Restricting Incoming Connections	719
Redirecting Users to Alternate Urls	721
Controlling Access to the HTTP Server	722
Throttling Client Access to Servers	725
Optimizing Access for Web and Mobile Clients	728
Configuring Access for OAB	728
Configuring Access for OWA	729
Configuring Access for Exchange ActiveSync	730
Configuring Access for ECP	731
Chapter 26. Optimizing Client Access Protocols	733
Managing RPC and MAPI over HTTP	734
Working with RPC and MAPI over HTTP	734
Configuring URLs and Authentication	736
Enabling the POP3 and IMAP4 Services	739
Optimizing POP3 and IMAP4 Settings	742
Configuring POP3 and IMAP4 Bindings	742
Configuring POP3 and IMAP4 Authentication	744
Configuring Connection Settings for POP3 and IMAP4	745
Configuring Message Retrieval Settings for POP3 and IMAP4	746
Chapter 27. Configuring Mobile Messaging	749
Mastering Mobile and Wireless Access Essentials	750
Getting Started with Exchange ActiveSync	750
Managing ActiveSync and OWA for Devices	750

Configuring Autodiscover	752
Understanding Autodiscover	752
Maintaining Autodiscover	752
Using Direct Push	757
Using Remote Device Wipe	758
Remotely Wiping a Device	758
Reviewing the Remote Wipe Status	761
Using Password Recovery	762
Recovering a Device Password	762
Managing File Access and Document Viewing	764
Configuring Direct File Access	764
Configuring Remote File Access	767
Integrating Office Web Apps Servers	768
Working with Mobile Devices and Device Policies	771
Viewing Existing Mobile Device Mailbox Policies	771
Creating Mobile Device Mailbox Policies	773
Optimizing Mobile Device Mailbox Policies	777
Assigning Mobile Device Mailbox Policies	779
Removing Mobile Device Mailbox Policies	781
Managing Device Access	783
Blocking Device Access	783
Using Access Rules	783
Setting Access Levels and Blocking Thresholds	785
Chapter 28. Tracking and Logging Exchange Server 2016	786
Configuring Message Tracking	787
Changing the Logging Location	787
Setting Logging Options	788
Searching the Tracking Logs	790
Beginning an Automated Search	790
Reviewing Logs Manually	791
Searching the Delivery Status Reports	794
Configuring Protocol Logging	796
Enabling or Disabling Protocol Logging	796
Setting Other Protocol Logging Options	796
Managing Protocol Logging	799

Optimizing Protocol Logging for HTTP	800
Working with HTTP Protocol Logs	801
Using Connectivity Logging	803
Configuring Connectivity Logging	803
Working with Connectivity Logs	804
Chapter 29. Maintaining Exchange Server 2016	806
Monitoring Events, Services, Servers, and Resource Usage	807
Viewing Events	807
Managing Essential Services	809
Monitoring Messaging Components	811
Using Performance Alerting	814
Tracking Memory Usage	814
Tracking CPU Utilization	816
Tracking Disk Usage	817
Working with Queues	819
Understanding Exchange Queues	819
Accessing the Queue Viewer	821
Managing Queues	823
Understanding Queue Summaries and Queue States	823
Refreshing The Queue View	824
Working with Messages In Queues	825
Forcing Connections to Queues	826
Suspending and Resuming Queues	826
Deleting Messages from Queues	827
Chapter 30. Troubleshooting Exchange Server 2016	828
Troubleshooting Essentials	829
Tracking Server Health	829
Tracking User and Workload Throttling	832
Tracking Configuration Changes	834
Testing Service Health, Mail Flow, Replication and More	835
Diagnosing and Resolving Problems	839
Identifying Recovery Actions	839
Identifying Responders	841
Identifying Monitors	842
Identifying Probes	844
Viewing Error Messages for Probes	845

Tracing Probe Errors	847
Troubleshooting Outlook Web App	851
Checking OWA Health	851
Understanding Unhealthy Status	852
Correcting Unhealthy Status	853
Index	2
About the Author	855